UCL

On the Resilience of the Dark Net Market Ecosystem to Law Enforcement Intervention

by

Cerys Bradley

PhD Thesis

Crime and Security Science

August 2019

Declaration of Authorship

I, Cerys Bradley, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed: _____

Date: _____

UCL

Abstract

Dark Net Markets (DNMs) are websites found on the Dark Net that facilitate the anonymous trade of illegal items such as drugs and weapons. Despite repeated law enforcement interventions on DNMs, the ecosystem has continued to grow since the first DNM, Silk Road, in 2011. This research project investigates the resilience of the ecosystem and tries to understand which characteristics allow it to evade law enforcement.

This thesis is comprised of three studies. The first uses a dataset contained publicly available, scraped data from 34 DNMs to quantitatively measure the impact of a large scale law enforcement operation, Operation Onymous, on the vendor population. This impact is compared to the impact of the closure of the DNM Evolution in an exit scam. For both events, the impact on different vendor populations (for example those who are directly affected and those who aren't) are compared and the characteristics that make vendors resilient to each event are investigated.

In the second study, a dataset acquired from the server the DNM Silk Road 2.0 is used to better understand the relationships between buyers and vendors. Network analysis and statistical techniques are used to explore when buyers trade and who with. This dataset is also used to measure the impact of a hack on Silk Road 2.0 on its population.

In the final study, discussions from the forum site Reddit were used to qualitatively assess user perceptions of two law enforcement interventions. These interventions were distinct in nature - one, Operation Hyperion, involved warning users and arresting individuals and the second, Operation Bayonet, actively closed a DNM. Grounded Theory was used to identify topics of conversation and directly compare the opinions held by users on each intervention.

These studies were used to evaluate hypotheses incorporated into two models of resilience. One model focuses on individual users and one on the ecosystem as a whole. The models were then used to discuss current law enforcement approaches on combating DNMs and how they might be improved.

Impact Statement

This research focuses on Dark Net Markets (DNMs), websites that facilitate the sale of drugs. This is a key crime area especially as DNMs have been used to trade Fentanyl, a priority drug for law enforcement in the UK. Therefore, this work has contributed to its research discipline but also to to current approaches to tackling cybercrime. The key impacts of each of the three studies including in this thesis are described below.

In the first study of this thesis, several methodologies for data preparation and validation within the study of DNMs were developed. In particular, this work presents a new technique for validating a publicly available dataset that has been used in multiple studies in this field. This is the first attempt to formally validate the dataset and determine what can reasonably used for research. The discussion of the dataset has implications for research already using the dataset and future research on datasets collected using the same methodology.

In order to conduct the second study in this thesis, a dataset was acquired from a law enforcement agency. This dataset gives a new insight on how buyers behave on DNMs. Buyers are an unstudied group because their activities are often hidden and so analysis of this dataset reveals new insights into the behaviour of these users. The results of this study have been used to comment on existing work using less complete datasets and contribute new findings.

The third study in this thesis presents a qualitative analysis of two law enforcement interventions. This is the first work to assess the impact of either intervention and so provides new insights into how they were received by the DNM ecosystem. It uses qualitative techniques which are rare within this discipline and so provides a different perspective, for example by revealing how individuals perceive the harms of law enforcement interventions on DNMs. The value of this work has been recognised through its acceptance at a workshop at the IEEE European Symposium on Security and Privacy, 2019.

Part of this research has been conducted in consultation with a law enforcement agency who provided data for this research. The results of this research are framed specifically for this agency and other law enforcement groups currently investigating DNMs. Several suggestions are made on how to improve the efficacy of law enforcement interventions on DNMs.

Acknowledgements

I would like to thank my supervisor Gianluca Stringhini for your patience, guidance and support. In addition, I would like to thank the UCL Crime and Security Science department for facilitating my PhD. I would especially like to thank Amy Clemens for answering an inordinate amount of inane questions.

Thanks to Rob Clarke, Samantha Dowling and the whole OCST R& A team for a much needed break from my research. You were so welcoming and I learned an enormous amount working with you. Thanks also to Steve Welsh and his team for my first internship and continued support of my research.

Thank you to Celine West and the UCL Museums for providing me years of enjoyable Saturdays talking about UCL research.

I owe a huge amount to Steve Cross for providing a creative outlet during my PhD and teaching me both how to share my research and why I should. Most importantly, thank you for introducing me to a fantastic community and helping me create my PhD support network.

Thank you to the Gym, Bloomsbury, especially to Andy and Angie, for providing a productive way to procrastinate for my thesis.

Enormous thanks to Anna, Oz, and Alex for supporting my projects, constantly checking in on me and providing the emotional support I have needed to complete this thesis.

I would like to thank my parents for supporting and encouraging me. As much as I have hated being asked when I will be finished every week for the past 6 months, I truly appreciate having you as my safety net. And, finally, I would like to thank Amanda for making sure I got through this in one piece.

Contents

Declaration of Authorship i

Abstract ii

Impact Statement ii

Acknowledgements iv

List of Figures ix

List of Tables xii

Abbreviations xv

1 Introduction 1
- 1.1 Dark Net Markets and Dark Net Technologies . . . 2
  - 1.1.1 Accessing a DNM . . . 2
  - 1.1.2 Creating an Account . . . 3
  - 1.1.3 Acquiring Funds . . . 3
    - 1.1.3.1 Multisig Transactions . . . 5
  - 1.1.4 Making a Purchase . . . 5
  - 1.1.5 Product Shipment . . . 6
  - 1.1.6 Feedback . . . 7
  - 1.1.7 Forums . . . 7
- 1.2 Who Uses DNMs? . . . 7
- 1.3 History of Law Enforcement on the Dark Net . . . 9
- 1.4 Contributions . . . 11

2 Literature Review 13
- 2.1 Dark Net Markets . . . 13
  - 2.1.1 Marketplace Studies . . . 14
    - 2.1.1.1 Collecting Data . . . 14
    - 2.1.1.2 Validating Data . . . 15
    - 2.1.1.3 Linking Vendors . . . 16

2.1.1.4 Categorising Products . . . 19
2.1.1.5 Calculating Sales . . . 20
2.1.2 User Behaviour . . . 21
2.1.2.1 Interview and Survey Studies . . . 21
2.1.2.2 Forums . . . 24
2.1.2.3 DNM Analysis . . . 28
2.1.3 Law Enforcement and the Dark Net . . . 33
2.2 Related Contexts . . . 37
2.2.1 Legal Highs . . . 37
2.2.2 Malware, Carding, and Cybercrime as a Service . . . 38
2.2.3 Hackers . . . 40
2.3 Resilience Theory . . . 43
2.3.1 Psychology . . . 44
2.3.1.1 Resilience in Individuals . . . 44
2.3.1.2 Community Resilience . . . 47
2.3.2 Ecology and Sociology . . . 47
2.3.3 Engineering . . . 51

3 Research Question and Hypotheses 54
3.1 Resilience Definition . . . 54
3.2 Adverse Events . . . 56
3.2.1 Impact on Ecosystem . . . 58
3.2.1.1 Ecosystem Hypotheses . . . 60
3.2.2 Impact on Users . . . 61
3.2.2.1 User Hypotheses . . . 64

4 Data 68
4.1 Gwern Branwen Dataset . . . 68
4.1.1 Information Extraction . . . 69
4.1.2 Data Validation . . . 70
4.1.2.1 Summary . . . 76
4.1.2.2 Review Data . . . 77
4.1.2.3 Ecosystem Coverage . . . 78
4.1.3 Forums . . . 80
4.2 Silk Road 2.0 Dataset . . . 81
4.3 Reddit Data . . . 82

5 Methodology 85
5.1 Cross Market Study . . . 85
5.1.1 Linking Vendor Accounts Across DNMs . . . 85
5.1.2 Measuring Vendor Lifetimes . . . 89
5.1.3 Collecting Product Information . . . 90
5.1.4 Measuring Vendor Reputation . . . 92
5.1.5 Identifying Events . . . 94
5.1.6 Continuing to Trade . . . 95
5.1.7 Vendor Movement . . . 95
5.2 Silk Road 2.0 Study . . . 95

5.2.1 Creating the Network . . . 96
5.2.2 User Information . . . 96
5.2.3 Network Analysis . . . 99
5.2.4 ERGM Analysis . . . 100
5.3 Reddit Study . . . 101
5.3.1 Measuring the Community . . . 101
5.3.2 Measuring the Impact of Events . . . 101
5.3.2.1 Quantitative Approach . . . 102
5.3.2.2 Qualitative Approach . . . 102

6 Results of Cross Market Study 103
6.1 General Statistics . . . 103
6.1.1 Vendors . . . 103
6.1.2 DNMs . . . 105
6.2 Analysis of the Impact of Events . . . 106
6.2.1 Operation Onymous . . . 106
6.2.2 Evolution Exit Scam . . . 110
6.2.3 DNM Closure . . . 114
6.2.4 Vendor Movement . . . 115
6.2.5 Summary . . . 116

7 Silk Road 2.0 Results 118
7.1 General Statistics . . . 118
7.1.1 Network Analysis . . . 119
7.1.2 UK Only Network Analysis . . . 122
7.2 Temporal Analysis . . . 124
7.2.1 How Does the Population Respond to Events? . . . 124
7.2.2 How Does the Shape of the Network Change Over Time? . . . 129
7.3 Vendor Buyer Relationships . . . 130
7.3.1 What Makes Vendors Popular? . . . 130
7.3.2 Can Vendor Behaviour Influence Buyers to Leave the Market? . . . 134
7.3.3 Bad Transaction Experiences . . . 138

8 Reddit Forums Results 144
8.1 General Statistics . . . 144
8.1.1 /r/DarkNetMarkets . . . 145
8.1.1.1 Network Analysis . . . 148
8.1.2 /r/dnmuk . . . 150
8.1.3 /r/Ebay . . . 151
8.1.4 Subreddit Comparisons . . . 152
8.2 Evaluation of Law Enforcement Interventions . . . 154
8.2.1 Activity . . . 154
8.2.2 Content Analysis . . . 158
8.2.2.1 Operation Hyperion . . . 158
8.2.2.2 Operation Bayonet . . . 169

9 Discussion 203
9.1 Hypotheses . . . 204

9.1.1 User Resilience . . . 204
9.1.2 Ecosystem Resilience . . . 209
9.2 Limitations . . . 213
9.2.1 Cross Market Study . . . 213
9.2.2 Silk Road 2.0 Study . . . 215
9.2.3 Reddit Forums Study . . . 216
9.3 New Findings and Future Work . . . 217
9.4 Recommendations for Law Enforcement . . . 218

10 Conclusion **222**
10.1 Cross Market Study . . . 223
10.2 Silk Road 2.0 Study . . . 225
10.3 Reddit Forum Study . . . 226
10.4 Findings . . . 228

A Validation Metrics **231**

B Username Matches **237**

C List of Subcategories **240**

D List of Common Words **243**

E Reputation **245**

F Timeline **250**

G CDF Plots of Advertised and Observed Listing Values **258**

H ERGM Results **263**

I ERGM Results **279**

Bibliography **281**

List of Figures

3.1 **Hypothesis Model for the Impacts of Events on Ecosystem** . . . 60
3.2 **Model for the Impacts of Events on Users** . . . 64

4.1 **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Vendors as well as the CDF Plot of the Smoothed Vendor Population. . . . 75
4.2 **Estimated Proportion of Vendors and Listings** Captured for All DNMs in Dataset. . . . 79
4.3 **Number of missing files** plot showing the number of missing listing files (green) and vendor files (pink) against the vendor population (purple). . 79
4.4 **Estimated Ecosystem Caputer** expected number of DNMs (blue) vs Observed Number of DNMS (green). . . . 80

6.1 **Number of Vendor Accounts** This plot shows the number of vendors controlling each number of vendor accounts. . . . 104
6.2 **Histogram of the Account Ages** This plot shows the ages of vendor accounts measured in days. . . . 104
6.3 **Ecosystem Population** Population change over time. It shows the number of new accounts from new vendors (purple), the number of new accounts from existing vendors (orange) and the number of existing accounts (green). . . . 105
6.4 **CDF of DNM Sizes** . . . 105
6.5 **CDF plot of the ages of the vendors** who were active during Operation Onymous. . . . 109
6.6 **CDF plot of the ranks of the vendors** who were active during Operation Onymous. . . . 109
6.7 **CDF plot of the number of accounts** owned by the vendors who were active during Operation Onymous. . . . 109
6.8 **CDF plot of the ages of the vendors** who were active during the Evolution Exit Scam. . . . 113
6.9 **CDF plot of the ranks of the vendors** who were active during the Evolution Exit Scam. . . . 113
6.10 **CDF plot of the number of accounts** owned by the vendors who were active during the Evolution Exit Scam. . . . 113
6.11 **Movement in Ecosystem Post Operation Onymous** This plot shows the DNMs active before and after Operation Onymous represented as nodes such that the size of the node is proportional to the number of vendors active on the DNM and the movement of vendors after Operation Onymous represented by directed edges. . . . 117

7.1 **Image of Network Coded by Type of Node.** This figure shows the network with buyer nodes coloured purple, vendor nodes coloured red, users that were both vendors and buyers coloured blue, and isolate nodes coloured green. . . . 119
7.2 **Image of Network.** This figure shows the network with nodes that have full profiles in the dataset coloured red and those that do not coloured green. . . . 119
7.3 **Population** This plot shows the number of users, buyers, and vendors in each week. . . . 124
7.4 **Number of Transactions** This plot shows the number of transactions that take place each week. . . . 125
7.5 **Amount of Sales** This plot shows the amount, measured in USD, of sales each week. . . . 125
7.6 **Amount of Sales and Change in Population** This plot shows the amount, measured in USD, of sales each week and the population over time plotted against each other for comparison. . . . 126
7.7 **Cumulative Population** This plot shows the cumulative population of all users (all), buyers (buyers), and vendors (vendors) against the population (population) over time. . . . 126
7.8 **Population Fluctuations** This plot shows the change in population each day. . . . 128
7.9 **Change in Population After Hack** This plot shows change in population and the change in existing population (users who were already on the DNM) immediately after the hack. . . . 128
7.10 **Density** This plot shows the density of the network over time. . . . 129
7.11 **Centralisation Over Time** This plot shows the indegree and outdegree centralisation over time. . . . 130
7.12 **Probability of an edge** being formed given a one unit increase in each of the variables. . . . 133
7.13 **Review Quality and Lifetime** This plot shows relationship between the proportion of negative reviews a vendor receives and the length of their lifetime. . . . 139
7.14 **Transaction Quality and Lifetime** This plot shows relationship between the proportion of negative reviews a buyer gives and the length of their lifetime. . . . 140

8.1 **Number of Users vs Number of Comments** this figure shows the number of comments and number of users for each post. . . . 146
8.2 **Score vs Number of Comments** this figure shows the score of each post plotted against the number of users. . . . 146
8.3 **Proportion of Contributor Type** this plot shows the proportion of contributors who made posts, comments or both during the measurement period. . . . 147
8.4 **Density** this figure shows the change in network density over time for subreddit */r/DarkNetMarkets*. . . . 149
8.5 **Centralisation** this figure shows the change in outdegree, indegree, and overall centralisation over time for subreddit */r/DarkNetMarkets*. . . . 149

8.6 **Reciprocity** this figure shows the change in network reciprocity over time for subreddit */r/DarkNetMarkets* as a range between the measure on the network formed with the minimum and maximum number of possible edges. 149
8.7 **Centralisation** this figure shows the change in transitivity over time for subreddit */r/DarkNetMarkets* as a range between the measure on the network formed with the minimum and maximum number of possible edges. 149
8.8 **Number of Users vs Number of Comments** this figure shows the number of comments and number of users for each post. 150
8.9 **Number of Contributors, Posts and Comments** in the DNM threads over the measurement period 155
8.10 **Number of Contributors and Moving Average** with 15 Week Window. 155
8.11 **Number of Posts and Moving Average** with 15 Week Window. . . . 156
8.12 **Number of Comments and Moving Average** with 15 Week Window. 156
8.13 **Number of Contributors and Moving Average** with 18 Week Window. 157
8.14 **Number of Posts and Moving Average** with 18 Week Window. . . . 157
8.15 **Number of Comments and Moving Average** with 18 Week Window. 158
8.16 **Keywords used to identify relevant posts** 159
8.17 **Number of Mentions of Each Theory** Over Time. 172
8.18 **Number of Comments** Containing Each Perspective on if Alphabay was Down Permanently or Temporarily 174
8.19 **Relative Proportions of Each Topic over Time** 180

9.1 **Model for the Impacts of Events on Users** (solid lines represent supported hypotheses, dashed lines represent partially support hypotheses and dotted lines represent unsupported hypotheses) 210
9.2 **Hypothesis Model for the Impacts of Events on Ecosystem** (solid lines represent supported hypotheses, dashed lines represent partially support hypotheses and dotted lines represent unsupported hypotheses, red has been used to highlight two hypotheses discussing the same relationship) 213

G.1 **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population. 259
G.2 **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population. 260
G.3 **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population. 261
G.4 **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population. 262

List of Tables

3.1 Specific Impacts on Ecosystem Attributed to Each Possible Event 59
3.2 Specific Impacts on Users Attributed to Each Possible Event 63

4.1 Sizes of Datasets 71
4.2 Kolmogorov-Smirnov Test Results on the CDF of the Advertised and Recorded Numbers of Listings 73
4.3 Kolmogorov-Smirnov Test Results on the CDF of the Number of Vendors Found in the Grams and Branwen Datasets 75
4.4 Review Data Available on DNMs 77
4.5 Comparison of the Estimated Proportion of Cloud 9 Vendors Included in Scrape 79

5.1 Category Classifier Accuracy 92
5.2 Location of Other Users 99

6.1 Chi Square Statistics 107
6.2 Average Values for Vendor Characteristics Before and After Operation Onymous. 107
6.3 The Variance Inflation Factor of each variable in the logistic regression model. 108
6.4 The Influence of the Variables in the Model Built to Predict if a Vendor Will Leave After Operation Onymous. 109
6.5 The Variance Inflation Factor of each variable in the logistic regression model. 112
6.6 The Influence of the Variables in the Model Built to Predict if a Vendor Will Leave After the Evolution Exit Scam. 112
6.7 Comparison of Vendors who Continued to Trade After the Evolution Exit Scam and Operation Onymous. 114
6.8 DNMs from the Observation Period that Closed Amicably 114
6.9 Results of ERGM analysis evaluating the variables of the size and age of DNMs in the ecosystem on the number of new DNM accounts made post Operation Onymous. 116

7.1 General Network Statistics 123
7.2 Influence of Vendor Characteristics on the Probability an Edge is Formed 132
7.3 Number of Weeks Each Variable Held Each Level of Significance 133
7.4 Measures of Variation for the Statistically Significant Values of the Results of ERGM Analysis Over Time 133
7.5 Logit Regression Results for the Model to Predict if a Buyer will Stop Trading when a Vendor Leaves the Market 136

7.6 Accuracy of the Model to Predict if a Buyer will Stop Trading when a Vendor Leaves the Market . . . 136
7.7 The Variance Inflation Factor of Each Variable in the Logistic Regression Model. . . . 137
7.8 The Distribution of Ratings Across All Reviews . . . 139
7.9 Accuracy of the Model . . . 140
7.10 Logit Regression Results of the Model to Predict if a Buyer will Stop Trading after a Bad Transaction Experience . . . 141
7.11 The Variance Inflation Factor of Each Variable in the Logistic Regression Model. . . . 142
7.12 The Variance Inflation Factor of Each Variable in the Logistic Regression Model. . . . 142
7.13 Logit Regression Results of the Model to Predict if a Buyer will Stop Trading 1 Day after a Bad Transaction Experience . . . 143

8.1 Summary of Subreddit Characteristics . . . 153

9.1 A Summary of the Conclusions on Each Hypothesis . . . 212

H.1 ERGM Results for Network Snapshot on 11/11/2013 . . . 263
H.2 ERGM Results for Network Snapshot on 18/11/2013 . . . 263
H.3 ERGM Results for Network Snapshot on 25/11/2013 . . . 264
H.4 ERGM Results for Network Snapshot on 02/12/2013 . . . 264
H.5 ERGM Results for Network Snapshot on 09/12/2013 . . . 264
H.6 ERGM Results for Network Snapshot on 16/12/2013 . . . 265
H.7 ERGM Results for Network Snapshot on 23/12/2013 . . . 265
H.8 ERGM Results for Network Snapshot on 30/12/2013 . . . 265
H.9 ERGM Results for Network Snapshot on 06/01/2014 . . . 266
H.10 ERGM Results for Network Snapshot on 13/01/2014 . . . 266
H.11 ERGM Results for Network Snapshot on 20/01/2014 . . . 266
H.12 ERGM Results for Network Snapshot on 27/01/2014 . . . 267
H.13 ERGM Results for Network Snapshot on 03/02/2014 . . . 267
H.14 ERGM Results for Network Snapshot on 10/02/2014 . . . 267
H.15 ERGM Results for Network Snapshot on 17/02/2014 . . . 268
H.16 ERGM Results for Network Snapshot on 24/02/2014 . . . 268
H.17 ERGM Results for Network Snapshot on 03/03/2014 . . . 268
H.18 ERGM Results for Network Snapshot on 10/03/2014 . . . 269
H.19 ERGM Results for Network Snapshot on 17/03/2014 . . . 269
H.20 ERGM Results for Network Snapshot on 24/03/2014 . . . 269
H.21 ERGM Results for Network Snapshot on 31/03/2014 . . . 270
H.22 ERGM Results for Network Snapshot on 07/04/2014 . . . 270
H.23 ERGM Results for Network Snapshot on 14/04/2014 . . . 270
H.24 ERGM Results for Network Snapshot on 21/04/2014 . . . 271
H.25 ERGM Results for Network Snapshot on 28/04/2014 . . . 271
H.26 ERGM Results for Network Snapshot on 05/05/2014 . . . 271
H.27 ERGM Results for Network Snapshot on 12/05/2014 . . . 272
H.28 ERGM Results for Network Snapshot on 19/05/2014 . . . 272
H.29 ERGM Results for Network Snapshot on 26/05/2014 . . . 272

H.30 ERGM Results for Network Snapshot on 02/06/2014 273
H.31 ERGM Results for Network Snapshot on 09/06/2014 273
H.32 ERGM Results for Network Snapshot on 16/06/2014 273
H.33 ERGM Results for Network Snapshot on 23/06/2014 274
H.34 ERGM Results for Network Snapshot on 30/06/2014 274
H.35 ERGM Results for Network Snapshot on 07/07/2014 274
H.36 ERGM Results for Network Snapshot on 14/07/2014 275
H.37 ERGM Results for Network Snapshot on 21/07/2014 275
H.38 ERGM Results for Network Snapshot on 28/07/2014 275
H.39 ERGM Results for Network Snapshot on 04/08/2014 276
H.40 ERGM Results for Network Snapshot on 11/08/2014 276
H.41 ERGM Results for Network Snapshot on 18/08/2014 276
H.42 ERGM Results for Network Snapshot on 25/08/2014 277
H.43 ERGM Results for Network Snapshot on 01/09/2014 277
H.44 ERGM Results for Network Snapshot on 08/09/2014 277
H.45 ERGM Results for Network Snapshot on 15/09/2014 278
H.46 ERGM Results for Network Snapshot on 22/09/2014 278
H.47 ERGM Results for Network Snapshot on 29/09/2014 278

I.1 ERGM Results for Network Snapshot on 06/10/2014 279
I.2 ERGM Results for Network Snapshot on 13/10/2014 279
I.3 ERGM Results for Network Snapshot on 20/10/2014 280
I.4 ERGM Results for Network Snapshot on 27/10/2014 280
I.5 ERGM Results for Network Snapshot on 03/11/2014 280

Abbreviations

BTC	**B**it**C**oin
CDF	**C**umulative **D**istribution **F**unction
Coin	**C**ryptocurrency
DNMs	**D**ark **N**et **M**arkets
ERGM	**E**xponential **R**andom **G**raph **M**odel
FRAM	**F**unctional **R**esonance **A**ccident **M**odel
FVEY	**F**ive Eyes Law Enforcement Group
LCS	**L**owest **C**ommon **S**ubstring
NP	**N**eyman **P**earson
NPS	**N**ew **P**sychoactive **S**ubstances
RFE	**R**ecursive **F**eature **E**limination
SVM	**S**upport **V**ector **M**achine
TF-IDF	**T**erm **F**requency-**I**nverse **D**ocument **F**requency
Tor	**T**he **O**nion **R**outer

Chapter 1

Introduction

Dark Net Markets (DNMs) are websites that facilitate the trade of illegal items. The most prolifically sold products are drugs (Christin (2013)), but many other items are available including hacking software, stolen goods, weapons, eBooks, and drug paraphernalia. They function in a similar manner to eBay or the Amazon marketplace by allowing buyers to create accounts and connect with independent sellers across the world. However, the use of Dark Net technologies allows users to anonymously exchange goods or services for cryptocurrencies such as Bitcoin.

The first DNM, *Silk Road*, was launched in 2011, and it has been estimated that the site had a sales revenue of $15 million in early 2012 (Christin (2013)), making approximately $1.2 million a year in commission (Soska and Christin (2015)). These values increased as the marketplace grew rapidly in size (Soska and Christin (2015)). This brought it to the attention of law enforcement and *Silk Road* was closed in 2013 by the FBI. Despite its closure and the arrest and conviction of the site administrator Ross Ulbricht (U.S. Attorney's Office (2015)), *Silk Road* was soon replaced by new DNMs (Soska and Christin (2015)). Since its closure, the ecosystem has continued to grow, surviving multiple, international law enforcement efforts. There have been over a hundred different DNMs in at least 5 different languages and the dominant DNMs now dwarf the original *Silk Road* with thousands of active vendors at a time, instead of in a lifetime (Soska and Christin (2015)).

Research on DNMs has identified the types of products being bought and sold, and given insight into the types of people buying and selling on DNMs as well as their motivations. Some studies have investigated the impact of law enforcement interventions on the DNM ecosystem, however existing research is mostly context specific with little comparison between different law enforcement approaches. This thesis examines the resilience of DNMs from individual users to the ecosystem itself in the face of multiple events

and answers the research question *Is the DNM ecosystem resilient to law enforcement interventions?*.

This chapter will describe DNMs, the technologies they rely on, the people who use them, and the law enforcement interventions that have attempted to combat them. It will then outline the results of this research.

1.1 Dark Net Markets and Dark Net Technologies

Silk Road, and its successors, provided their users with the ability to buy and sell products with relative anonymity and financial safety. To do this, they relied on a number of different technologies, including the Dark Net, cryptocurrencies, and PGP encryption. These are each explained below through the process of making a purchase on a DNM which followed the *Silk Road* escrow model.

1.1.1 Accessing a DNM

In order to participate on a DNM, users must first have access to the Dark Net[1]. This is a collection of networks within the Internet that host websites which are not indexed by search engines[2] and have restricted access either by being invitation only or having specific technological requirements for users. The dominant network is the Tor Network, a technology created by the U.S. Navy, which is reliant on a global network of volunteers (Tor Project (2017)). Websites hosted on Tor have the domain `.onion`.

Silk Road was only accessible through Tor, as are most of the DNMs, however a few have also, or exclusively, been accessible through another anonymous network, I2P. The Tor Network, and I2P, hide the location of the servers hosting websites on their network and protect users from traffic analysis.

In the Tor or I2P network, instead of communicating directly, a device will connect to a server through a randomly selected path of other devices in the network. When many people are using the network, the different paths intersect with each other obscuring the links between any one device and the server they are visiting. This can make it difficult,

[1] This term has different meanings in different contexts, for example it can also reference a "network telescope" which allows a large-scale event on the Internet to be observed. The Dark Net, in the context of this work, is also commonly referred to as the Dark Web or the Anonymous Web however the term Dark Net has been chosen in this dissertation because it is directly referenced in the term Dark Net Market.

[2] It has been pointed out by an examiner that, in a sense, DNMs are indexed if they have been collected by Google via onion.to.

though not impossible, for law enforcement to determine who is visiting which websites and who is hosting them.

In practice, accessing the Dark Net hosted by I2P or Tor involves downloading free, open source software. Both provide web browsers with user interfaces similar to, for example, Google Chrome or Mozilla Firefox. Because Dark Websites are not indexed by search engines, users must know the precise address of the site they want to visit, however many surface websites, such as `DeepDotWeb.com`, provide the addresses of different DNMs.

1.1.2 Creating an Account

When a DNM has been identified, users must register an account. This process does not require personal information or identity verification, applicants simply have to create a username and password and set a PIN to withdraw money (Christin (2013)). Vendors pay a small fee which could be as little as $2 or as high as several hundred dollars and then create a public profile on which they can advertise their products. Buyers are often able to sign up for free and do not have a profile page or public information to maintain. Some sites are invitation only and, as the ecosystem has grown, vendors have begun to advertise invitation links for popular DNMs as products.

Vendor profiles display different types of information on different DNMs. Some allow vendors to display only their description, reputation rating, and products. Others allow any combination of the following: the country they are shipping from, the number of sales they have completed, the reviews they or their products have received, when they began trading and were last logged onto the site, any other verified accounts they have active on other DNMs, and, the number of sales they currently have in progress.

1.1.3 Acquiring Funds

In addition to membership, in order to make purchases, buyers need access to a cryptocurrency. *Silk Road* exclusively used Bitcoin, which is the dominant currency within the ecosystem, however alternatives employed by other DNMs included Litecoin, Dogecoin, Darkcoin and Monero. Cryptocurrencies are online currencies built around a cryptographic protocol that allows them to be spent and received. Each of these currencies is decentralised, meaning that transactions are approved through a community of users, instead of a designated authority. This feature means that the transactions cannot be regulated by law enforcement or similar actors e.g., transactions cannot be blocked or accounts closed in the same way that PayPal, for instance, will freeze accounts they suspect are linked to illegal activity.

The manner in which transactions are approved is dependent on the currency. Here the Bitcoin protocol will be briefly outlined as this is the most commonly used cryptocurrency in the ecosystem.

All Bitcoin transactions in the cryptocurrency's history are recorded on a central blockchain which is completely public. In order to approve a new transaction, it must be added to the blockchain and, once a transaction is recorded on the blockchain and has been built upon by other blocks in the chain, the transfer of funds is considered complete.

Transactions are added to the blockchain in blocks by members of the Bitcoin community. Anyone can add a block to the blockchain and they are rewarded for doing so with Bitcoin, this process is often referred to as *mining* as the Bitcoin they are rewarded are new and created specifically for this purpose. In order to add a block, several conditions must be satisfied:

- the block must be valid, i.e. it has not broken any of the rules that would facilitate users double-spending their Bitcoin or otherwise receiving Bitcoin they have not earned;
- the block must contain a value that produces a specific output when hashed. The value can only be found through a brute force method that is designed to be computationally expensive. This is the proof-of-work protocol that Bitcoin relies on to keep its users honest (Nakamoto (2008)). The protocol asks users to expend a large number of resources that will only be recuperated if their block is added to the blockchain, i.e. that they also follow the first condition;
- the final condition is that the majority of the community accept the block, this is determined by other users building on the block to further the chain. If the block is in the longest chain on the blockchain then it is considered a valid block.

Whilst participating in the creation of Bitcoin requires technical knowledge, as well as a considerable amount of computing power, as with Tor and I2P, the everyday use of Bitcoin is much simpler than the technologies they rely on. Bitcoin can be purchased and traded online using open source apps which mimic online banking interfaces. To date, there are nearly 16.7 million Bitcoin in circulation (Blockchain Luxembourg S.A. (2017)), worth approximately $95 million (Coindesk (2017)). However, the prices can fluctuate dramatically (Ciaian et al. (2016)).

Because the blockchain is public, every Bitcoin transaction is public, and it is therefore possible to trace every coin back to when it was first created. However, creating a wallet or address to receive payments does not always require identity verification. As such, individuals who are able to entirely separate their wallet from their real world identity can use Bitcoin to make anonymous purchases online. Further, it is possible for users to

generate new Public Keys (the identifiers of their Bitcoin wallets) for each new purchase thus making it difficult for observers to prove that their transactions are linked.

When a buyer has acquired cryptocurrency, they must deposit funds into their account. This process is not usually instantaneous and, as many DNMs will not allow you to initiate a purchase without sufficient funds, making a purchase can often take a lot longer than on a surface web shopping site. Any funds in a user's wallet, be they the profit of a sale or intended for a purchase, are held by the administrators of the DNM. This means that users must trust the DNM to stay active so that they can access their funds and that the admin will approve any withdrawals they want to make.

1.1.3.1 Multisig Transactions

More recent DNMs, for example the now closed *Hansa Marketplace*, facilitate multi-signature (multisig) transactions. To conduct a multisig transaction, the buyer transfers funds to a cryptocurrency wallet which has been constructed with one public key but multiple private keys. The funds can then be transferred when the transaction is signed by a predetermined number of the private keys.

For example, in a 2 of 3 multisig transaction, 3 users (the buyer, the vendor and the marketplace) all have unique private keys but, in order to make a transfer, only 2 must sign it. On a DNM, this means the buyer and the DNM can cooperate to approve a refund, the vendor and buyer can cooperate to finalise a sale, and the vendor and buyer can cooperate to finalise a sale or refund even if the DNM has gone offline, making users less vulnerable to the harms of exit scams.

1.1.4 Making a Purchase

DNMs provide two types of listings, public listings which are advertised on the site and stealth listings that are only visible with an invitation. Stealth listings are usually only made available to buyers that the vendor trusts and has conducted business with previously. Not all advertised products are available – a common practice for vendors who have run out of stock on a particular item is to inflate the price beyond an amount anyone would be willing to pay (Christin (2013)). This allows them to retain the item's reviews for when they are able to continue selling it. If buyers are unable to find a product they want to purchase, many DNMs offer a messaging service which allows buyers to approach vendors and request custom listings more suited to their interest.

Products can be purchased two ways – via escrow or finalising early (F.E.). In an escrow purchase, the buyer transfers the funds to the DNM itself, who will then hold them until

the buyer confirms they have received the product (or "finalises" the purchase). If the product does not arrive and the buyer is unable to resolve the issue and guarantee suitable compensation from the vendor, their funds are returned, minus the escrow fee charged by the DNM. When a buyer agrees to F.E., they immediately mark the purchase as complete, even though their product has not arrived. This option is usually insisted upon by vendors selling to new buyers as a way of protecting themselves from scammers who will claim their product has not arrived when it had been sent in order to avoid paying. Some DNMs do not allow F.E. purchases as it enables vendors to scam buyers by receiving payment and not shipping goods.

1.1.5 Product Shipment

If a physical product was purchased, the vendor must then ship it to the buyer. In order to do this they need the buyer's address. DNMs provide a secure messaging service which allows buyers to communicate their addresses, though it is often recommended that buyers and vendors use PGP encryption as well.

PGP encryption is an encryption protocol which allows a sender and receiver to communicate such that their intercepted messages cannot be decrypted unless any eavesdropper has access to the cryptographic secret used to encrypt the message. For a sender to send an encrypted message to a receiver, the receiver must first generate two keys – a public key and a private key. The public key can be made public and is communicated to the sender, they use this to encrypt their message and, once that is completed, the message can only be decrypted by the private key which is known only to the receiver.

Whilst open source technology also exists for PGP encryption, this is the most technical of the three technologies described and the hardest to master. It has been observed that PGP adoption on DNMs has fluctuated over time and, in particular, increased after large law enforcement interventions (Soska and Christin (2015)).

When a vendor has an address they must package and send their product. Vendors use a variety of different techniques to transport drugs, and other items, across borders or countries without arousing suspicion. For instance, drugs are often vacuum sealed to hide their smell and packaged in business style envelopes so that large numbers of items posted regularly from the same address do not look suspicious (Van Hout and Bingham (2014)). Vendors are rated on their stealth, as well as the quality of their products and timeliness of their arrival. Additionally vendors try to ship packages at different times and from different locations aware that suspicious and regular postal activity can be investigated by law enforcement.

Buyers are discouraged from using their real addresses, instead they may direct shipments to neighbours' houses or PO boxes, or abandoned buildings and then intercept or collect the parcel. They may also use fake names in order to be able to deny a package containing drugs is intended for them, however, vendors will likely refuse to ship using obviously fake names for fear of arousing suspicion from postal inspectors.

1.1.6 Feedback

On *Silk Road*, when a purchase was finalised, it was mandatory for buyers to complete a feedback form on the vendor or product. Not all DNMs require feedback, but many do. It has been found that the overwhelming majority of feedback on DNMs is positive (Soska and Christin (2015)), though this is not necessarily a reflection of all purchases being a positive experience. A similar phenomenon has been observed on the auction site eBay (Resnick et al. (2006)). Some vendors will require buyers to leave positive feedback in order to get compensation if their product has not arrived. When the purchase is F.E., buyers are unable to leave informative feedback but may return to edit their feedback if they have had a particularly poor or positive experience.

1.1.7 Forums

In addition to reviewing each other on the DNM, users may share their experiences on DNM related forums. These forums can be hosted by the DNM itself or might be a separate Dark Website or, even, a thread on a surface web site such as `Reddit.com` (Reddit User (2017)). They are used to review different vendors and buyers, discuss different drug types and their safe usage (Caudevilla (2016)), advertise products; discuss law enforcement operations (Lacson and Jones (2016)), evaluate new DNMs and, debate the reason existing DNMs have closed.

1.2 Who Uses DNMs?

Interviews and surveys on users of DNMs, particularly *Silk Road*, have presented a profile of young professionals and students, predominately male, with varying previous experience buying or selling drugs offline (Van Hout and Bingham (2013b, 2014)). For example some buyers, interviewed in 2013, had a drug history of 18 months and others 25 years (Van Hout and Bingham (2013b)). The most popular products have, consistently, been drugs and, more specifically, cannabis, MDMA, and stimulants (Soska and Christin (2015)). Users were technically proficient in the technologies required for DNM use,

however many found *Silk Road* out of curiosity or from reading news stories, as opposed to actively seeking an online environment on which to buy and sell drugs (Van Hout and Bingham (2013b, 2014)).

The trade on DNMs appears to be concentrated to a small number of countries. Dittus et al. (2017) found that 70% of trade for cannabis, cocaine and opiates was concentrated to the U.S., U.K. Australia, Germany, and the Netherlands (Dittus et al. (2017)). These countries were found to be some of the most popular destinations and locations in other studies (Christin (2013)) however they are not countries associated with the production of many drugs. This has inspired the conclusion that vendors are primarily located in consumer countries (as opposed to producer countries) and that DNMs are an additional mechanism for the sale at the consumer end as opposed to a facilitator of drug trafficking (Dittus et al. (2017)).

One of the primary, reported motivations for purchasing drugs on *Silk Road*, or other DNMs, was the convenience of being able to access a large range of drugs and have them delivered directly to you (Phelps and Watt (2014)). This, combined with the reassurance provided through the review based system and accompanying user forums (Phelps and Watt (2014); Van Hout and Bingham (2013b,a); Barratt et al. (2014); Christin (2013)) makes purchasing drugs online as reliable (if not more so (Phelps and Watt (2014))) as in person but with significantly less effort requirements. A number of users have also reported that the experience is safer because, if you are not meeting face to face, there is not the same opportunity for violence (Ormsby (2016)). Further, the escrow and resolution systems available in *Silk Road*, made users feel more confident in participating in deals (Van Hout and Bingham (2013b)). Finally, the perceived anonymity was also described as a crucial element that attracted users to *Silk Road* (Barratt et al. (2014)) as it made users believe that there was reduced risk of being caught buying or selling drugs.

However, users of *Silk Road* have also described being wary of the reliability of reviews as an indicator of good service (Van Hout and Bingham (2013b)) and described gaining confidence in the system only after several successful transactions (Van Hout and Bingham (2013a)), thus indicating that poor user experience could be a deterrent to using DNMs. Despite the measures put in place by *Silk Road* administrators, some users were still the victims of scams (Phelps and Watt (2014)) which caused inconvenience even if they did not leave the customer out of pocket. In addition, DNMs have closed down or exit scammed stealing thousands of dollars from customers in the process with potential implications on the trust placed in *Silk Road* (Soska and Christin (2015)).

By analysing responses to the Global Drugs Survey (2014), Barratt et al. (2016) compared the experiences of DNM users when buying online to when they purchased from

friends or other drug dealers. They found that users reported less instances of threats to personal safety, concerns over drug impurities, and concerns of law enforcement involvement when buying online than in either of the other two scenarios. However they reported more instances of loss of money and products not being received when buying on DNMs (Barratt et al. (2016)).

The discussed studies, that present findings from interviews or surveys of self selected respondents, imply that there are many benefits to trading in online, as opposed to offline, DNMs. However, there are still some risks associated with making online purchases that do have the potential to be exploited in order to make trading online seem less appetising.

Analysis conducted by Munksgaard and Demant (2016) on the forums of several DNMs was used to measure the evolution of community values over time. In this study, the topics of forums are analysed to determine the main discourse in the community. This study showed that dominant topics were business orientated (e.g. about the distribution or consumption of products) and that a topic expressing libertarian discourse rose in popularity between 2011 and 2013 before becoming less popular after the closure of *Silk Road* (Munksgaard and Demant (2016)). This finding, along with the interviews and surveys on the motivations and opinions of users potentially points to an economics-minded user base that is invested in the ecosystem for its profitability.

1.3 History of Law Enforcement on the Dark Net

There have been five major law enforcement efforts on the ecosystem: the closure of *Silk Road*, Operation Commodore, Operation Onymous, Operation Hyperion, and Operation Bayonet and the closure of *Hansa*.

Silk Road was closed in October 2013 by the FBI (Van Hout and Bingham (2014)). The FBI used several different strategies to close the site not all of which have been officially disclosed, these included infiltrating *Silk Road* as vendors and having direct conversations with the site admin Ross Ulbricht whilst undercover (Zetter (2013)). When *Silk Road* was closed, Ross Ulbricht was arrested and has since been convicted of seven narcotics and money laundering charges. He received a sentence of life imprisonment in April, 2015 (U.S. Attoney's Office (2015)). At the time of the closure, the FBI reportedly seized $3.6 million in Bitcoin from *Silk Road* (Clark (2013)) as well as the servers themselves.

Several studies have investigated the impact of *Silk Road's* closure on the ecosystem. A measurement of the ecosystem over time shows that the volume of sales was reduced immediately after the closure of *Silk Road*, but this rapidly grew again, surpassing the

height of the volume of sales observed on *Silk Road* within 6 months of its closure (Soska and Christin (2015)). Further, there is evidence to suggest that the event did not deter others from creating new DNMs or using them (Aldridge and Décary-Hétu (2015); Décary-Hétu and Giommoni (2016); Lacson and Jones (2016); Munksgaard et al. (2016); Soska and Christin (2015); Van Buskirk et al. (2015)) . Potential reasons for why the intervention did not appear to stop the growth of the ecosystem include the idea that the media stories that accompanied *Silk Road's* closure may have advertised it to new customers and that the FBI's actions closed a large monopoly allowing other DNMs to grow (Buxton and Bingham (2015)).

In February, 2014 Dutch Police shut down the DNM Utopia in Operation Commodore. The operation resulted in 5 arrests and the seizure of BTC900, or $610,900 (DeepDotWeb (2014f)). As with the closure of Silk Road, Dutch Police gathered evidence by going undercover onto the DNM, throughout the operation they made multiple purchases and were hired for an assassination (DeepDotWeb (2014f)).

Operation Onymous, the largest law enforcement intervention on the ecosystem, occurred in November, 2014. The FBI, Europol, and the Department of Homeland Security conducted a joint investigation on Tor sites (Greenberg (2014b)). They seized 414 .onion sites, of which 11 were DNMs: *Silk Road 2.0*, *Bluesky*, *Torbazaar*, *Cloud9*, *Topix2*, *Hydra*, *Alpaca*, *Cannabis Road 3*, *Flugsvamp*, *Black Market*, and *Pandora* (Greenberg (2014b)). In addition, 17 arrests were made including the administrator of *Silk Road 2.0*, Blake Benthall (U.S. Attorney's Office (2014b)) and his deputy Brian Farrell who has since been sentenced to 8 years in prison (U.S. District Court for the Western District of Washington at Seattle (2017)). Similar to the closure of *Silk Road*, an initial reduction in activity was observed however academic research on the event has concluded that any impact the intervention has would be short term (Décary-Hétu and Giommoni (2016)).

The fourth intervention, Operation Hyperion, was of a different nature to its predecessors. It took place in November, 2016 and whilst it was also an international effort, this time involving the Five Eyes Law Enforcement Group (FVEY), an intelligence alliance between Australia, New Zealand, Canada, the UK and the US, it did not result in the closure of active DNMs. Instead, law enforcement groups in different countries approached individuals who were thought to have been active on the ecosystem and warned them about the potential of arrest should they fail to cease their activities (FBI (2016)). Swedish Police have claimed to have spoken to 3,000 suspects, New Zealand Police have stated they approached 160, and the FBI 150 (Drugs Forum User 5-HT2A (2016)). Dutch Police created a Dark Net Site and publicly named DNM users under investigation (Harfenist and Turgeman (2016)). As of yet, no academic research has been published assessing the impact of this intervention.

The fifth and most recent intervention took place in July, 2017. This was a semi-coordinated effort between the FBI, the DEA, and the Dutch Police which resulted in the closure of two DNMs *Hansa* and *Alphabay*. Dutch Police seized *Hansa* in June 2017, however instead of shutting the DNM, they kept it active to collect information on users (Europol (2017)). They then approached the FBI and DEA who were coordinating Operation Bayonet which was comprised of a series of raids across Canada and the seizure and shut down of *Alphabay* on 4 July 2017 (Europol (2017)). The closure of *Alphabay* was, at first, disguised as an exit scam as this was predicted to deter more users than a law enforcement intervention (Europol (2017)). *Hansa* was kept running as Dutch Police correctly predicted the displaced *Alphabay* population would move to this DNM after its closure allowing them to collect information on a greater proportion of the population. A DNM, *Dream Market*, that was not affected by the intervention, has continued to see an increase in its population since Operation Bayonet (BBC (2017)).

1.4 Contributions

This thesis presents three studies that evaluate multiple law enforcement operations both from a quantitative and qualitative perspective. A public dataset containing data collected from over 80 DNMs is used to compare the impact of Operation Onymous to the Evolution Exit Scam on the vendor population. An as of yet unresearched dataset taken from the server of *Silk Road 2.0* is used to understand vendor and buyer interactions and measure the impact of a hack on the user population and site trade. Posts and comments from relevant Reddit forums are used to qualitatively understand and compare the repercussions of Operations Hyperion and Bayonet and the closure of *Hansa*.

Through this research, this thesis makes several contributions to the literature, both in the form of new methodologies and in the form of new knowledge. These contributions are as follows:

- a new methodology for comparing and linking public vendor profiles using profile descriptions and product listings which is shown to identify more vendor pairings than existing approaches;
- a new methodology for comparing vendor reputations across different DNMs allowing for the impact of cross market events to be measured on this variable;
- a comparison of Operation Onymous and the Evolution Exit Scam which demonstrates how the law enforcement operation has a greater impact on the vendor population than the Evolution Exit Scam but that this impact is limited to vendors operating on affected DNMs;

- an adaptation on several network analysis measures, such as density, that give a more precise measurement of the interactivity of the *Silk Road 2.0* marketplace;
- a measurement of the hack of Silk Road 2.0 which is shown to have little permanent effect on the size of the population but a temporary reduction in trade;
- analysis of vendor buyer relationships that demonstrates vendor age, reputation, and diversity have marginal positive impacts on the probability a vendor will conduct trade, though each of these variables is shown to have a smaller influence than in other, similar studies;
- logistic regression analysis of buyer decisions to leave *Silk Road 2.0* when their trading partner had left the DNM showing that buyers were more likely to leave if they had just experienced a bad transaction or had a positive relationship with the vendor who had left;
- a qualitative study of two Reddit forums */r/darknetmarkets* and */r/dnmuk* which compared Operation Hyperion and Operation Bayonet and the closure of *Hansa* showing that the latter led to more posts and comments and appears to have had a greater impact on the user population.

Chapter 2

Literature Review

This section outlines the literature pertaining to the study of DNMs in the fields of Information Security and Criminology as well as some similar contexts such as websites selling "legal highs" and underground forums. It also presents relevant literature on Resilience Theory from Psychology, Ecology, and Engineering.

2.1 Dark Net Markets

A number of exploratory studies have been conducted to understand the nature and scope of DNMs beginning with (Christin (2013)) which attempted to quantify the amount of money passing through *Silk Road*. In order to conduct this research, techniques for collecting data had to be developed and methodology devised to make claims about the number of active participants and sales on each DNM. These studies have looked at the evolution on the overall ecosystem and its response to law enforcement efforts as well as the motivations and behavioural patterns of individual users.

The purpose of DNMs and manner in which they are hosted mean that they pose a difficult research challenge, particularly when collecting data. As a consequence, the research presented in this section is often limited either by a small sample size or incomplete picture. Further, many variables, such as the total revenue of a site, must be measured indirectly. It is perhaps for these reasons that there is not a complete consensus of the exact nature of DNMs and their sustainability.

Whilst the existing literature has explored many research questions, such as evaluating the size of the ecosystem, what products are sold on DNMs and from where, and how the ecosystem has reacted to certain law enforcement efforts, there are still a number of questions left unanswered. For instance, it is still unclear how the most recent law

enforcement efforts, Operations Hyperion and Bayonet, have impacted on DNMs. Most relevant to this work, Resilience Theory has yet to be applied to the ecosystem, an approach which may help to explain why DNMs respond to law enforcement in the way they do, i.e. with new DNMs being created or gaining dominance and the population and sales volumes continuing to grow despite initial, immediate reductions.

2.1.1 Marketplace Studies

2.1.1.1 Collecting Data

Many studies on DNMs have been conducted on data that has been scraped directly from the Dark Net (Broséus et al. (2016); Soska and Christin (2015); Christin (2013); Décary-Hétu and Giommoni (2016); Dolliver (2015a); Munksgaard et al. (2016); Duxbury and Haynie (2017); Nurmi et al. (2017)). Of these studies, some (Décary-Hétu and Giommoni (2016); Munksgaard et al. (2016); Broséus et al. (2017)) rely on a dataset collected by the independent researcher Gwern Branwen. This dataset contains scrapes of 89 DNMs collected between 2013 and 2015 (Branwen et al. (2015)). Each scrape is a partial snapshot of a marketplace and so does not contain every vendor profile or listing page on each day (Décary-Hétu and Giommoni (2016)), the scrapes were also not conducted at consistent time intervals across all markets, as the size of the DNM and its uptime (when it was active and not crashed), as well as potentially other unknown constraints, restricted how often scrapes could be conducted.

Whilst several studies have use this dataset, it is recognised as being severely limited. Conclusions made on analysis conducted with the data should be made cautiously and steps are required to clean and validate the data. For example, to overcome the irregularity of data collection, (Décary-Hétu and Giommoni (2016)) aggregated the data on a weekly basis to create more complete pictures of the DNMs studied.

Collecting data in this manner can lead to incomplete datasets for DNMs that enact anti scraping measures such as rate limiting, because the DNM has low availability, or because of Tor network issues that lead to pages not fully downloading (Van Buskirk et al. (2015)). Further, a poorly designed tool may not be able to follow every link on a DNM or might follow the wrong links and log itself out. This is why some researchers have developed tools specifically for this purpose, (Soska and Christin (2015); Branwen et al. (2015)) as opposed to using freely available tools such as HTTrack (Christin (2013); Dolliver (2015a)) which have been shown to be less capable through experience (Christin (2013)). If a tool crashes without registering an error it may not even be clear that it has failed to document large sections of the DNM, (Munksgaard et al. (2016)). An

alternative to this automated process is to manually download each page on a DNM which is limited by the labour intensity of this method, (Van Buskirk et al. (2015)).

Some studies attempted to collect one full snapshot of the DNM or DNMs being studied, i.e. a complete picture of all the vendors and products advertised on the site at the time the scrape was initiated, (Aldridge and Décary-Hétu (2014); Dolliver (2015a); Van Buskirk et al. (2015); Duxbury and Haynie (2017)) whereas others looked at multiple scrapes over an extended time period, (Broséus et al. (2016); Christin (2013); Décary-Hétu and Giommoni (2016); Soska and Christin (2015); Munksgaard et al. (2016); Wadsworth et al. (2017)). When looking at multiple DNM studies, some researchers have chosen to select only a few DNMs, instead of all of those active at the time of the study. For example, Soska and Christin (2015) only included DNMs that had mandatory feedback policies and transaction volumes greater than $1,000 (determined by inspection). Décary-Hétu and Giommoni (2016) looked at the top 5 DNMs measured by their number of listings with the belief they were representative of the entire ecosystem during the measurement window. Others, limited their sample size by country, (Broséus et al. (2016)).

Not all studies have focused exclusively on the Dark Net. Wadsworth et al. (2017) also collected data from the surface Web. To do so they searched for relevant terms on Google and Bing included the top five search results in their analysis. The search terms were "buy" plus one of the following key phrases "legal highs", "research chemicals", "bath salts", "party pills", and "herbal highs". This was repeated throughout October 2015. This was compared to the New Psychoative Substances (NPS) listing pages and the profiles of vendors selling NPS products downloaded from 22 DNMs for two days in October, December 2015 and February, April, June, August, and October 2016, (Wadsworth et al. (2017)).

2.1.1.2 Validating Data

In order to assess the completeness of the data analysed, a number of different techniques were employed. To verify the completeness of their data, Dolliver and Kenney (2016) compared the number of listings available on the site to the number they found in their dataset. This was similar to the approach taken by Dolliver (2015b), who manually inspected the DNM scraped in their study in order to verify their tool was collecting the expected amount of information. They found many fewer products that were actually available to buy than the number advertised though their approach has been criticised (Dolliver (2015b)) as the results it produced are contradictory to other, similar studies.

This technique has implied that DNMs can falsely advertise that a greater number of products are available than is the case.

Soska and Christin (2015) used the Schnabel estimator to quantify the amount of feedback omitted from their web-scrapes. The Schnabel estimator allows for a population size to be approximated by the method of repeat sampling when said sampling is conducted uniformly at random from a constant population, (Schnabel (1938)). Scrapes within a 60 day period were considered to be samples of the same population, the Schnabel estimator was then employed to approximate the proportion of information that was captured across different numbers of scrapes. This approach led to the conclusion that 10 or more independent scrapes could capture 90% of a DNM, (Soska and Christin (2015)). However, as the Schnabel Estimator assumes a constant population (meaning that, at least, the ration of sampled and unsampled members remains constant even if the population size changes) and the size of DNMs can be be volatile, this method is not necessarily appropriate for measuring the full population.

In addition to this approach, a post analysis method was also used to validate the data and methodologies. Ground truths from publicly released information about law enforcement investigations into Ross Ulbricht and Blake Benthall (such as the sales volumes of the relevant DNMs) in combination with leaked seller pages from the DNM *Agora* were compared to the analytical results. These spot checks produced very similar findings to the study implying the techniques used produced accurate estimates of the amount of money passing through the sites analysed, (Soska and Christin (2015)).

2.1.1.3 Linking Vendors

After the fall of *Silk Road*, the overall ecosystem became much more volatile with DNMs operating for shorter periods of time and disappearing without warning. In 2014, the Dark Net search engine Grams was launched allowing users to find vendors and listings across multiple DNMs, (Reddit User (2014c)). These two developments are perhaps why vendors began operating on more than one DNM at a time from 2014, (Soska and Christin (2015)). Indeed, Soska and Christin (2015) found vendors operating on up to 6 DNMs between 2013 and 2015. In order to study vendor behaviour, researchers therefore needed to develop a way of linking different accounts across different DNMs. A failure to do so can also lead to an over representation of the size of the overall vendor population. For example, Soska and Christin (2015) collected a dataset of 29,528 unique usernames which they reduced to a presumed 9,386 unique vendors.

The simplest method of connecting accounts is by matching usernames. Vendors operating on more than one marketplace may choose to be easily identifiable in order to

leverage their reputation from other DNMs, (Soska and Christin (2015)). Another motivation may be to better retain their customer base if a DNM they are operating on is closed down. Broséus et al. (2016) chose to match case when matching username strings, finding 146 unique vendors from a pool of 198, whereas Soska and Christin (2015) did not. In either case, stolen or impersonated usernames will result in misidentification using this method (Broséus et al. (2016)), and connections between usernames that differ by punctuation will be missed, (Soska and Christin (2015)). As such, both studies also compared PGP keys as two accounts controlled by different vendors would not advertise the same PGP key, (Broséus et al. (2016); Soska and Christin (2015)). Broséus et al. (2016) showed that most, but not all, accounts linked by the same PGP had similar usernames which at most differed by only the case or the addition of a word. However, this metric is also imperfect as not all vendors generated their own PGP key and, as keys can expire, some vendors must use more than one key, (Broséus et al. (2016)).

Van Buskirk et al. (2017) compared vendor names after removing ASCII characters that were not numbers or letters and spaces. They also removed common suffixes such as 'the'. Vendor matches were inspected for duplicates based on common words (e.g. 'drugs' and 'therealdrugs') which were not considered aliases. This approach reduced the vendor population collected by 52.3% (compared to a 67.9% reduction by Soska and Christin (2015) who also utilised PGP key verification).

In order to link similar but not identical usernames, methods such as the Levenshtein distance (Décary-Hétu and Giommoni (2016)) or longest common substring (Iofciu et al. (2011)) can be used. However, these can capture common words, such as 'cannabis', and mistakenly link accounts. Décary-Hétu and Giommoni (2016) manually inspected profiles with usernames that had a Levenshtein similarity of less than 25% of the length of the username string. This approach is not necessarily feasible for large sets of vendors. In order to automate the process of inspecting profiles, other data, such as profile descriptions can be compared using, for example, Term Frequency-Inverse Document Frequency (TF-IDF) analysis (Iofciu et al. (2011)) or other tools that determine linguistic patterns in text.

Spitters et al. (2015) conducted authorship analysis on the forum for the DNM Black Market Reloaded. They used a forum to identify collections of accounts they suspected to be controlled by the same user despite operating under different usernames. Data was collected from October 2012 until December 2013 and consisted of 92,333 posts by 8,348 accounts in 12,923 threads.

A combination of topic independent features, n-grams, and time-features were extracted from posts at least 5 words in length. A Support Vector Machine (SVM) classifier was used to attribute posts to authors. In order to test the accuracy of the classifier,

a test dataset was created by simulating aliases for existing accounts and distributing the posts associated with the account between it and the fictional alias. The classifier achieved between 88 and 94% accuracy depending on the number of users it was trained on (Spitters et al. (2015)). The accuracy of any classifier built for this purpose will be heavily dependent on the amount of profile information created by users because this determines the amount of material the classifier can be trained on. As such, it may be easier to conduct this analysis in forums than on the sites themselves where there is comparatively less text created by each user.

In addition, Wang et al. (2018) have argued that stylometry based analysis can be unreliable because vendors do not have strong writing styles and that matches may be heavily dependent on whether vendors used the exact same sentences on their different accounts. Indeed they showed that vendor writing styles were not as strong as templates provided by the DNM, this, and the use of different languages, can make it difficult to match vendors using this technique (Wang et al. (2018)). Instead, Wang et al. (2018) compared the photographs posted by vendors and used deep neural networks to determine photography styles that could be used to identify Sybils.

Their method was evaluated on the Gwern Branwen dataset. A ground truth was created by artificially creating separate accounts from the same vendor and dividing that vendor's photos between the artificial accounts. When tested, their method was shown to have high accuracy and outperformed their stylometry based method (Wang et al. (2018)). When applied to the Branwen dataset (focusing just on *Agora*, *Silk Road 2.0*, and *Evolution*) the method found 850 pairs (from 8,691) which were then manually inspected using characteristics such as name spelling and meaning, profile contents and photos. 738 of these pairs were evaluated as a "confident yes", only 484 of which had similar usernames (Wang et al. (2018)).

This method is susceptible to vendors using publicly available photos, particularly those posted by other vendors. As such, the authors recommend that it is used to help reduce the resources required to match vendor accounts manually and could be paired with a human element to confirm matches.

A final method for linking vendor accounts across sites is to use the information provided by the DNMs themselves. Some DNMs, such as Black Market Reloaded, allow vendors to advertise their accounts on other DNMs. Further the search engine Grams has an InfoDesk feature which produces aggregate information on specific users, this was utilised in (Soska and Christin (2015)) to link vendors. This method is only applicable to those vendors who choose to provide this information to the DNMs they advertise on and those operating at the same time as Grams.

Whilst some vendors will be motivated to ensure their different accounts are linkable (Soska and Christin (2015); Broséus et al. (2016)), there is also an incentive for vendors with poor reputations to begin trading under a new alias in an untraceable manner. It is also plausible that vendors create multiple accounts on the same marketplace in order to carry out a Sybil Attack and give themselves positive feedback, or others negative feedback (Christin (2013)). As such, the number of vendors estimated as active on the site may always be an overestimate of the number of real world operations behind the usernames.

2.1.1.4 Categorising Products

Some studies have sought to understand the types of products sold. For example, to find evidence of vendors selling different products on different DNMs and make conclusions about their operation set up (Broséus et al. (2016)) or the overall market in their country (Broséus et al. (2017)). Further, looking at changes in the market share of different categories of drugs may reveal market responses to law enforcement efforts. Indeed, Soska and Christin (2015) found a significant increase in the market share of cannabis after four major events - two law enforcement actions, the closure of two DNMs and a hack. This may have been a result of vendors changing their practices in order to lessen the impact of being arrested or scammed (Soska and Christin (2015)), if so monitoring the market share of different categories would be a good indicator of the perceived risk associated with trading on DNMs.

However, categorising different products is not a trivial problem. Whilst the simplest solution would be to use the categories given by the DNMs themselves, many DNMs have different category lists and different definitions for each drug type making cross market studies difficult. For example, *Silk Road 2.0* categorised Benzodiazepines as a subcategory of "Prescriptions", whereas *Oxygen* considered them an independent category "Benzos". Additionally, there is no consensus within the literature on a list of drug categories and their definitions. There are different ways of defining drugs for example by choosing to use their chemical composition (Broséus et al. (2017)) or their effects (Nurmi et al. (2017)) and each produces a different set of categories (Lee and Antin (2011)).

Given a set of categories, different methods have been employed to sort listings into the correct category. Christin (2013) and Soska and Christin (2015) used the category information provided by the vendors. However, the reliability of this method has been called into question (Coopman et al. (2016); Van Buskirk et al. (2016a)) as vendors may mislabel their listings. Instead, Graczyk and Kinningham (2015) designed an algorithm

to classify product listings on Dark Net Markets into 12 categories (11 drug types and one category for other products) using TF-IDF analysis and an SVM classifier. They achieve 79% accuracy, when testing on a training dataset, claiming this is greater than other tools (Graczyk and Kinningham (2015)). They trained their algorithm on listings from the Gwern dataset. Soska and Christin (2015) used a similar approach, they were able to use 1,941,538 pre-labelled listings on *Evolution* and *Agora* as a ground truth and used tf-idf tokenisation as input to an L2-Penalized SVM under L2-Loss. They used 10 fold cross validation to evaluate their classifier as 98% accurate (Soska and Christin (2015)).

Van Buskirk et al. (2016a) sorted the drug listings into 12 distinct categories using the Excel 2010 lookup function. This categorised 79% of the listings and the remaining 21% were done by inspection creating an issue of scalability for larger studies, 250 could not be categorised and so were excluded from the study. The categories were informed by population-level substance surveys from the US, Australia, and the UK.

Durrett et al. (2017) consider this problem in the context of cybercrime forums and build a classifier that can categorise products across different forums. Posts in the forum were labelled at the token level to allow for more than one product per post, 478,176 of which were hand annotated by three annotators in order to train their classifiers. Whilst Neyman-Pearson (NP) prediction at the post level was found to be the most accurate classifier, it was still not as accurate as the human annotators (Durrett et al. (2017)).

Another characteristic of products that may want to be documented is their weight. To identify the weights of different products on the DNM *Silk Road*, Przepiorka et al. (2017) inspected each of the 6,126 products manually in order to extract their weights. No automated approaches to this problem could be found.

2.1.1.5 Calculating Sales

Most DNMs do not advertise the number of sales each vendor has made. As such, studies seeking to quantify the total revenue of sites have used the number of pieces of feedback as an approximator for this value (Christin (2013); Soska and Christin (2015); Aldridge and Décary-Hétu (2014).) It is presumed that one piece of feedback equals one sale (Aldridge and Décary-Hétu (2014)), however this is not necessarily the case if users buy multiple items in a single purchase (Soska and Christin (2015)). This method, at best, will provide an underestimate as not all the feedback can necessarily be collected in a scrape. In (Soska and Christin (2015)), it was estimated that only 60% of all feedback was collected and feedback does not capture stealth (or private) sales (Christin (2013)).

Further, this method is only applicable to DNMs that make feedback compulsory and provide highly granular feedback time stamps.

Nurmi et al. (2017) used the stock information (the number of products still available) on the product listing pages of the DNM Silkkitie. If this data is accurate, it is potentially a more reliable approach as it will capture the number of products sold even when a buyer makes multiple orders in the same purchase and it should also reflect the number of stealth sales. However, this information is not universally available and so can only be utilised on some DNMs, further no formal process of measuring the accuracy of the information has been presented in the literature. The information may be false if vendors are able to manually input the figures, and they may be incentivised to falsely lower the stock levels in order to encourage purchases before stock runs out. Further, technical difficulties in the site may present false information. However, these issues are likely comparable to measuring sales through review data as vendors are able to post false reviews from fake buyer accounts and the review data may also be presented inaccurately by the server.

2.1.2 User Behaviour

The anonymity provided by DNMs, as well as their illicit nature, can make studying user behaviour difficult. Users are anonymous and often times incentivised to publish misleading information to deceive either law enforcement or other DNM users. This problem is exacerbated when focusing on buyers - whereas vendors have profiles which list their ratings and reviews, and are required to pay fees to trade, often buyers are able to keep their profile anonymous, can join the site for free, and are not required (or able) to build a reputation. As such, it is often not possible to produce even basic quantitative estimates, such as how many unique buyers there are or their average lifetime. There are three main approaches to studying users in the existing literature: conducting interviews and surveys on users of DNMs, analysing DNM forums, and, analysing the publicly available information on the DNMs.

2.1.2.1 Interview and Survey Studies

To gain insight into buyer behaviour or motivations, some researchers have used interviews (Van Hout and Bingham (2013b,a, 2014); Bancroft and Reid (2016); Lavorgna (2016)) or made use of the Global Drugs Survey (Barratt et al. (2014)). Whilst this method provides an opportunity to gather information directly from users it can be limited due to small or skewed samples. Further, the majority of interview and survey

based studies have focused on users of the now closed site *Silk Road*, which may have provided a different user experience to other DNMs.

Van Hout and Bingham (2013b) conducted 20 interviews with *Silk Road* users recruited through the site. The interviews took the form of an open ended questionnaire sent through a secure messaging relay. The results were combined with 168 screen shots, 4 threads and 1,249 posts on the *Silk Road* forum. The study presented an image of the DNM being populated by educated, recreational drug users who favoured "MDMA, 2C-B, Mephedrone, nitrus oxide, ketamine, cannabis, and cocaine" (Van Hout and Bingham (2013b)). The participants cited both economic motivations (a wider range of drugs was available at better prices) and reasons of safety (both in terms of the quality of the product and the greater physical safety) as explanations for making online purchases (Van Hout and Bingham (2013b)). They sought out vendors based on trust and qualities such as speed of transaction, ability to provide stealth in delivery and product quality (Van Hout and Bingham (2013b)). In general, the anecdotes of using the site were positive with few examples of poor service or law enforcement experience (Van Hout and Bingham (2013b)). This was a follow up study to (Van Hout and Bingham (2013a)) in which one DNM user was interviewed.

Van Hout and Bingham (2014) looked at vendors. One researcher spent 12 months as an active participant on *Silk Road* forums in order to create a "rapport" with users, permission was granted by the administrator of the site and 10 vendors were interviewed via the secure messaging system (they were presented with open ended questions and responded via PGP encrypted messages). Vendors were attracted by the "low risk, high traffic, high mark-up" environment (Van Hout and Bingham (2014)). They described a culture of harm reduction and responsible vending as well as a dedication to providing quality service through professional advertising and communication, visibility on forums, speedy dispatch, stealthy and overweight packages, and competitive pricing (Van Hout and Bingham (2014)).

Given the small sample sizes and voluntary recruitment method, the results may provide a skewed perspective on most user experiences. It is difficult to assess whether or not these experiences are specific to the subset of DNM users who were both active at the time of the study and willing to talk to researchers or instead represent the typical experience (Lund Research Ltd (2012)). As such, Van Hout and Bingham (2013b) supplemented their results with data and quotes collected from the messaging boards in the *Silk Road* forum. A similar method was employed by Bancroft and Reid (2016) who conducted analysis on the forums of an anonymous DNM and used the results to inform 5 qualitative interviews. The results of the forum analysis corroborated the results of the analysis conducted on the interviews (Van Hout and Bingham (2013b)).

Recruiting a larger sample size can be difficult. Van Buskirk et al. (2016b) interviewed 745 participants for the 2014 Ecstasy and related Drugs Reporting System (EDRS) during March – July 2014. Participants were over the age of 16 and had used a psychoactive stimulant drug at least once a month in the six months preceding the study, they were recruited using a purposive method (i.e. chosen because they fit this criteria as opposed to at random (Teddlie and Yu (2007))). Of these participants, only 82 reported using the Dark Net and of these, only 60% had done so in the past year (Van Buskirk et al. (2016b)). The results of this study corroborated the work in (Van Hout and Bingham (2013a,b)).

Outside of academia, users of DNMs have also been interviewed by investigative journalists. Orsmy (2016) evaluated the results of these interviews providing new insights, for example that some users chose to share an account with multiple people in order to increase their reputation faster and be able to change addresses more regularly to evade law enforcement (Lavorgna (2016)). However, they are still not necessarily representative of the *Silk Road* user base (Lavorgna (2016)).

A potential method of expanding the size of a study sample and widening its reach is through a survey, as opposed to interviews. Barratt et al. (2014) used the Global Drugs Survey (GDS) which surveyed drug users from the U.S., Australia, and the U.K. about their experiences including making online purchases. The data was collected from November, 2012 to January, 2013 and returned responses from 9,470 respondents. 65% of USA, 53% of Australian, and 40% of UK respondents had heard of *Silk Road* and 18, 10 and 7% (respectively) had made purchases on *Silk Road* (Barratt et al. (2014)). (This compares to 13.2%, 8.2% and 25.3% of respondents of the 2017 GDS from the USA, UK, and Australia who stated they had made a purchase from a DNM in the previous 12 months.)

As with the interview participants, users cited wider range, better quality, greater convenience, and the vendor rating system as the most common reasons for buying on *Silk Road*. Adequate access to products and fear of being caught were the most common reasons for not using *Silk Road*. Logistic Regression was used to show that the appeal of *Silk Road* was dependent on "country-specific deterrents and market characteristics", however differences could also be explained by the varying demographics determined and methods of survey distribution (Barratt et al. (2014)). There were some limitations to this study also, as the survey was distributed through social media sites and a music magazine, limiting its scope. However, it has a much greater sample size of respondents and the Global Drugs Survey has accurately represented drugs trends in the past (Barratt et al. (2014)).

2.1.2.2 Forums

Many DNMs hosted an accompanying forum for users to discuss products, other users, the vulnerability of the DNM to law enforcement and other related topics. Forums can be a useful medium for understanding the motivations and opinions of DNM users. Indeed, Caudevilla (2016) cites forums as a useful way to gather information such as emerging drugs trends. This is because information collected on forums is potentially a more honest image of user behaviour if the participants are unaware they are being observed and can also allow for a greater sample size particularly on DNMs as there are many active forums. Quantitative analysis of forums can also potentially reveal information about how the community is shaped and how much time is invested into community discussion.

However, studies of this nature are limited in a similar way to those which look at the DNMs themselves. The anonymity of users makes it difficult to test the validity of the published material as users are able to create multiple accounts and spread false information, further the collection of data, particularly through scrapes, may result in an incomplete dataset (Munksgaard and Demant (2016)). Analysis of these forums has focused on user reactions to law enforcement interference (Lacson and Jones (2016)), user perceptions of drug quality and opinions on use (Caudevilla (2016); Bancroft and Reid (2016)), and the community structure of DNM users including their political ideologies.

In order to evaluate the impact of the FBI seizure of *Silk Road* on DNM users, Lacson and Jones (2016) coded comments and discussion threads from the forums of *Silk Road*, *Agora* and *Evolution* over the five-month period after the shut down of *Silk Road*. Comments were hand coded on their content, the type of user (their reputation and experience level), and the type of comment (whether it was deemed positive, negative, or neutral by the coders) in order to draw conclusions about the types of topics being discussed, who was participating in them and how. They found that more experienced users were more likely to seek an alternative to *Silk Road 2.0* and that comments coded as neutral by the researchers scored higher average karma scores within the forum whereas those comments coded as negative scored, on average, negative karma scores (Lacson and Jones (2016)).

Horton-Eddison and Di Cristofaro (2017) also used forums to understand how the closure of *Silk Road* affected community members. They focused specifically on attitudes towards the escrow system employed by both *Silk Road* and *Silk Road 2.0* and used Corpus Linguistics Assisted Discourse Analysis (CLADA) to extract opinions from the DNM's forums. For this study, CLADA involved taking an automated approach to

identify collocates (co-occurring terms) to the word escrow and locating statistically significant instances of the use of these collocates before qualitatively attributing sentiment to both the word escrow and its collocates (Horton-Eddison and Di Cristofaro (2017)). These instances were separated chronologically by the FBI closure of *Silk Road* and the hack of *Silk Road 2.0* so that the change in sentiment towards escrow services and shift towards decentralised systems, caused by these event, could be measured.

Caudevilla (2016) set up a number of threads in active DNM forums. These threads focused on the safe use of drugs and asked for questions on the theme of medical advice on the DNMs *Silk Road*, *Silk Road 2.0*, and *Evolution* (Caudevilla (2016)). Quantitative data, such as how many comments were posted and the total number of visitors were collected over time and showed an active engagement in the threads with the number of visits reaching the tens of thousands (Caudevilla (2016)). It was believed that the threads were responded to positively (Caudevilla (2016)).

Bancroft and Reid (2016) also observed forum data in order to assess user perceptions of drug quality and how their value judgements are made. From March 2015, they observed the forum of a large, competitive DNM (which they chose to anonymise), it had 152 threads which ranged in size from 20 to 7,000 comments created over a 2 year period. In May 2015, they collected 3196 posts. The comments were handcoded using a coding system that was adapted with new content until the coding system was exhausted by the content, i.e. all the content was coded. The researchers looked at the different context and definitions of words such as quality and purity as well as discussions surrounding dosage and use and cooking and manufacture.

The coding exercise informed questions created for 5 interviews in the summer of 2015. The study concluded that DNM users also make street purchases, that assessments of quality are produced collaboratively, and that forums could facilitate the dissemination of information about drug safety (Bancroft and Reid (2016)).

To understand what was being discussed on DNM forums, Luo (2017) analysed the forums of the DNM *Agora* using the publicly available dataset collected by Gwern Branwen. They used TF-IDF analysis and Latent Dirichlet Analysis (LDA) to identify topics of discussion in the forum in order to understand its purpose and relation to drug trafficking. LDA is a Bayesian approach to classifying, clustering, and summarising text (or other discrete data) (Blei et al. (2003)). The approach assigns each item within a corpora a finite number of topics where each topic is defined by an infinite combination of probabilities (Blei et al. (2003)), as such it allows for items to be defined by more than one topic and for topics to overlap. This approach identified 10 topics and 50 threads were selected from each topic, totalling 2,578 posts, which were manually parsed and hand coded using Grounded Theory. Grounded Theory is a data driven approach to

determining the key issues concerning, or processes utilised by, the research subjects and is implemented by reading content to determine such issues/processes and then rereading to substantiate their existence within the community (Grounded Theories Ltd (2016)). Using this approach, the posts were coded into 4 themes: invitation to treat (posts relating to business activities), risk management and social control, drug-related knowledge exchange, and community support.

It was concluded that the main functions of the *Agora* forum were business facilitation, order maintenance and social support (Luo (2017)). It was argued that the forum aided the DNM *Agora* in the facilitation of drug trafficking as it created an environment in which participants had access to knowledge on how to commit the crime, were given the impression that "everyone was doing it" and were provided with evidence that other members of the community needed drugs thus lowering social inhibitions about participating in drug trafficking (Luo (2017)). These conclusions were made on the basis of qualitative discussions around case studies identified within the forum and the Crime Science literature. Based on the discussion within the study, they can be taken as hypotheses as opposed to findings.

Munksgaard and Demant (2016) relied on forum data to investigate the community values of DNM users. The market forums of *Silk Road*, *Silk Road 2.0*, *Evolution*, *Agora*, and *Black Market Reloaded* were scraped between October 2013 and March 2015 by the independent researcher Gwern Branwen (Munksgaard and Demant (2016)). Posts from English language subforums were extracted and processed to remove duplicates and reposts, these were then reduced to bags of words. Unsupervised Topic Modelling was used to extract the different topics discussed in each forum over time and their dominance. This was used to track the prevalence of libertarian ideals within the marketplace forums. Their analysis suggests that, whilst libertarianism was highly discussed at the launch of *Silk Road* and through its life, when *Silk Road* was taken offline the dominance of the libertarian discussion decreased dramatically and has never since risen back to the same level (Munksgaard and Demant (2016)). Further, no other, competing political discourses were found suggesting that libertarian discourse achieved a "hegemonic, dominant position" (Munksgaard and Demant (2016)). These results were taken, in the context of DNMs, to imply that, whilst the motivations of users may have been political in the early days of *Silk Road*, they had shifted to being more business focused (Munksgaard and Demant (2016)).

Lorenzo-Dus and Di Cristofaro (2018) used the forums of DNMs *Silk Road* and *Silk Road 2.0* to understand the concept of trust and its evolution in the community. They analysed a total of 24 forums across the two DNMs using scraped data collected by independent researcher Gwern Branwen (Lorenzo-Dus and Di Cristofaro (2018)). They

used a Corpus Assisted Discourse Studies (CADS) approach which, in this context, involved identifying key words associated either with trust or with figures within the ecosystem (e.g. "vendor") and using custom built Python scripts to extract relevant posts or words associated those concepts in order to analyse how they are discussed.

They found that vendors were discussed in a polarised manner, i.e. in an extremely positive light or a negative one and concluded that the forums were used to identify vendors who could and could not be trusted (Lorenzo-Dus and Di Cristofaro (2018)). When analysing the relationship between trust and the DNMs themselves, the authors used examples of posts to demonstrated common linguistic styles used by posters. For instance, posters were observed sharing intimate facts about themselves (such as physical descriptions needed for advice on drug taking) implying trust in the forum and demonstrating trustworthiness in themselves by referring to their relevant experience and authority in discussion topics (Lorenzo-Dus and Di Cristofaro (2018)).

A key finding was that the discussion of the centralised escrow format was more positive on the *Silk Road* forum whereas posters on the *Silk Road 2.0* forum appeared to have a preference for decentralised models (Lorenzo-Dus and Di Cristofaro (2018)). This claim is corroborated by the examples and case studies given in the paper, however quantitative evidence is only alluded to and not explicitly included in the paper.

In order to assess the methods of conflict resolution in DNMs, Morselli et al. (2017) conducted a study on 10 marketplaces *Alphabay*, *Dream Market*, *Valhalla*, *Hansa*, *Python*, *Acropolis*, *Tochka*, *Cryptomarket*, *Outlaw*, and *Nucleus*. The rules of engagement from each market as well as 200 discussion threads from the scam reports section of the forum of one DNM were analysed. Each post in the thread was hand coded as being resolved through one of the following: Tolerance, Avoidance, Ostracism, Third-party Intervention, Negotiation, or Threats. Each conflict was coded as being caused by Transaction Failures, Scams, Bad Market Management, Unfair Competitive Practices, Social Interactions, or Law Enforcement Activities. The majority of conflicts were found to be caused by transaction failures or scams and most conflicts were caused by nondelivery, lack of communication from the buyer, or low product quality (Morselli et al. (2017)). The conflict resolution strategies of ostracism (identifying a vendor with a scam), third-party mediation, tolerance, and avoidance were used most frequently, negotiation and threats used more scarcely and violence not at all (Morselli et al. (2017)). A common pattern of conflict resolution was ostracism, followed by third party mediation and then banishment. Another was ostracism followed by tolerance or avoidance, the former was usually an indicator of a successful resolution. The rules of engagement seem to be enforced by a "scamwatch" a dedicated group of community members. When threats were made, they were usually monetary threats or threats against the reputation of

the other party. This study was only conducted on public data, it could therefore only look at the conflict resolution strategies used after private messaging failed, implying that marketplaces are even better at conflict resolution, or it could omit the threats of violence which are more likely to take place in private, implying the opposite (Morselli et al. (2017)).

Rekšŋa et al. (2017) conducted quantitative analysis on DNM forums to understand if the users of the forums behave according to a specific structure or interact more randomly. The specific structure evaluated is the Configuration Model, a method for generating graphs which begins with a predetermined degree sequence (a set of predetermined values for the number of edges each vertex has in the network) which must sum to an even number and then randomly joins vertices until each has the correct number of edges for their assigned degree (Barabási et al. (2016)).

Networks were created from 26 forums by representing contributors in the forum as nodes and drawing edges between them if one contributor replied to another by embedding a post by the first contributor in a post of their own. Then, a number of graph features, including the degree of the nodes in the networks, the networks' assortativity, and how the networks are clustered, were measured for these networks and compared to predicted values calculated on random networks generated using the Configuration Model.

It was concluded that the Configuration Model could not accurately describe the networks of DNM forums (Rekšŋa et al. (2017)). In particular, the nodes of smaller degrees were clustered more in the observed networks than in the predictive model however more regular users did behave in a manner similar to that predicted by the Configuration Model (Rekšŋa et al. (2017)).

2.1.2.3 DNM Analysis

Finally, user behaviour has also been analysed by looking at information publicly available on the site (Aldridge and Décary-Hétu (2014); Hardy and Norgaard (2016); Van Buskirk et al. (2016a); Dolliver and Kenney (2016).) Whilst these observations can reveal patterns in user behaviour, they cannot necessarily provide conclusions on the motivations or reasoning of users. However, they do lead to and evidence interesting hypotheses, such as that buyers consider a sudden change in vendor behaviour as a more useful indication of a scam than consistently poor behaviour (Hardy and Norgaard (2016)). Observational studies have been used to investigate purchasing and selling behaviour (Christin (2013); Aldridge and Décary-Hétu (2014); Soska and Christin (2015); Dolliver and Kenney (2016),) the role of each vendor's reputation in a DNM's economy (Hardy

and Norgaard (2016); Nurmi et al. (2017); Przepiorka et al. (2017), relationships between vendor location and the products they sell (Van Buskirk et al. (2016a); Broséus et al. (2016, 2017); Dittus et al. (2017),) and the community structure of the ecosystem (Duxbury and Haynie (2017)).

Purchasing and Selling Behaviour. Christin (2013) conducted a study on the DNM *Silk Road* by scraping the site between February and July of 2012. 2,485 unique items were found across the sites 220 categories and 564 distinct sellers with at least one listing advertised were found. Variance in vendor population was observed around events such as pot day (20 April) for which a large number of vendors joined temporarily and when popular vendors went on hiatus. The majority of sellers remained on the marketplace for 100 days or less and only 112 stayed for the lifetime of the marketplace.

Feedback data was used to evaluate the level of customer satisfaction. 97.8% of the 184,809 pieces of feedback analysed were positive (the transaction was rated as 4 or 5 out of 5), a figure not too disimilar to eBay (Christin (2013)). Of the 20,844 pieces of feedback that signalled the buyer had had to finalise their transaction early, only 342 were retroactively edited to be negative, this was taken to mean that only a small proportion of buyers were scammed through the early finalisation feature with the caution that this number is likely an underestimate as the reaction requires additional effort from the reviewer (Christin (2013)).

Aldridge and Décary-Hétu (2014) conducted a study on *Silk Road* to examine the claim that the DNM was an "eBay for drugs" i.e., mostly used by recreational drug users making purchases for personal use. They scraped the site from 13 – 15 September 2013 and collected approximately 13,000 pages, 1,084 vendor pages and 11,904 active listings. In order to identify which listings were purchased for personal use and which for resale, they first eliminated 51 listings with inflated prices designed to signal the item was out of stock. The remaining listings were ordered by product price and then divided into five quintiles. In order to find the quantity of each listing, a variable presented differently, and sometimes not at all, in different listings, a random sample was taken from each quintile and the quantities were handcoded for the listings in each sample. In the top two quintiles, the mean purchase price for each drug category was in the thousands of dollars, these products were purchased less frequently, and the average unit price was closer to the estimate of the street value of each drug than those in lower quintiles. This led to the conclusion that the top quintiles were items for resale.

Soska and Christin (2015) analysed 16 DNMs from November 2011 until June 2015. They looked at active vendors (vendors with at least one active listing) and found that a large increase in vendor population since the closure of *Silk Road* (Soska and Christin (2015)). More than 10% of sellers were active throughout the whole measurement period

though half were present for only 220 days or less (Soska and Christin (2015)). About 70% of vendors did not sell more than $1000 of product and only 2% were able to sell more than $100,000. Sales volumes were calculated as the product of an item's price and the number of reviews it received. Analysis on the adoption of PGP keys showed that, whilst on *Silk Road* between 2/3 and 3/4 of vendors user PGP keys, this reduced after the market was closed in late 2013 and then rose after Operation Onymous potentially due to vendors feeling a need for greater privacy protection (Soska and Christin (2015)).

Dolliver and Kenney (2016) crawled *Agora* in June 2015 and *Evolution* in September 2015. They found 2,325 unique vendor profiles and 43,382 listings, 29,147 of which were for controlled substances. Logistic regression, t-tests, and chi-square tests were used to identify statistically significant differences in behaviour between vendors who sold drug related products and those who did not. Some differences were found, for instance in the average price of drug and non-drug related items (Dolliver and Kenney (2016)).

Reputation. Hardy and Norgaard (2016) investigated whether or not a vendor's reputation allowed for premium prices by comparing the quality of the reviews different vendors received and their prices for similar products. This technique has also been applied to eBay (Resnick et al. (2006); Melnik and Alm (2002); McDonald and Slawson (2002)). They concluded that, on *Silk Road*, a fall in the reputation of a vendor had a bigger impact on the price than the vendor's reputation overall. They used sales data on 9,604 sales from *Silk Road 2.0* from November 2013 until October 2014. The sales were of 119 cannabis listings from 41 vendors and the transaction volume ranged from 1 sale to 668 sales with each of the listings weighing 15 grams or less in order to maximise the chances are for personal use (this weight was chosen as it is below the felony possession amounts in Florida and Virginia). The data was collected using a scrape, of the 30,000 pages downloaded 300 were not collected due to a scrape error and the errors were all concluded to be random. Prices were converted to dollars to standardised for fluctuations in BTC and were used to approximate the trust a buyer placed in the vendor, under the assumption that buyers pay a premium for reliable vendors. The logistic regression tests conducted on the results show that the item rating is more important than the vendor rating and that a fall in a vendors rating is considered a more worrying factor than the overall rating of the vendor. The results also showed that buying larger weights resulted in better value (Hardy and Norgaard (2016)).

A similar study was conducted by Nurmi et al. (2017) on the Finnish DNM Silkkitie Nurmi et al. (2017). Analysis on 260 vendors and 3823 product listings collected using daily scrapes between November 2014 and September 2015 found that both vendor reputation and a vendor's product availability (number and diversity of products) affected a vendor's sales (Nurmi et al. (2017)). In addition, this study found that vendors are

active for only a short period, 62.8 days, and that many products are never sold (Nurmi et al. (2017)).

Przepiorka et al. (2017) also investigated the relationship between vendor reputation and their ability to trade, this time on *Silk Road*. They compared vendor reputations to the prices and popularity of their products. They found that vendors increased their prices if their reputation increased and decreased prices when their reputation fell (Przepiorka et al. (2017)).

Location. Van Buskirk et al. (2016a) looked at the locations of vendors on *Agora*. They quasi-randomly selected 7 days between the 15 March and 15 February then manually downloaded the drug category pages on each day. It was shown that the number of listings differed from country to country, as did the types of drugs. Further, the distribution of drug categories was country specific, e.g. vendors based in the U.S. sold disporportionately more cannabis in comparison to other countries (Van Buskirk et al. (2016a)).

Broséus et al. (2016) conducted a study on 8 different DNMs (*Agora*, *Blue Sky*, *Evolution*, *Silk Road 2.0*, *Cloud 9*, *Pandora*, *Hydra* and *Andromeda*) by scraping data between August and September of 2014. They collected 198 vendor profiles and 3,685 listings belonging to Canada based vendors. They found that most vendors operate on one site and have less than 10 listings and the users who operate on more than one site tend to have more listings (Broséus et al. (2016)). They also found that vendors sell different products on different sites, potentially as a reaction to market demand (Broséus et al. (2016)).

Broséus et al. (2017) investigated trends in the shipping destinations of products available on the DNM *Evolution*. They collected information on 4,171 vendors and 92980 listings from the scrape conducted by Branwen. Their analysis showed that the marketplace is dominated by English-speaking and Western countries and different countries are at the forefront of different product categories (China sells the most NPS, India the most Prescription Drugs, the Netherlands and Canada the most prescription and illicit drugs). Further, some countries exhibit different trends in shipping destinations, for instance the U.S and Australia ship more domestically whereas the Netherlands, Germany, and China ship more internationally.

In order to understand the impact of DNMs on the global drug trafficking network, Dittus et al. (2017) looked at the advertised locations of cannabis, opiate, and cocaine vendors on 4 large DNMs active in June and July of 2017. They used data on product seizures and self reported drug use collected by the UNODC World Drug Report to identify producer and consumer countries, respectively. These were compared to the

number of trades on the DNMs (as measured by the number of reviews) and their associated geographical data (the country each vendor claimed to be operating from). A disparity between the countries that produce drugs and the locations of the vendors potentially implies that the trade facilitated by DNMs is concentrated to that which happens in consumer countries as opposed to the drug trafficking network as a whole (Dittus et al. (2017)).

Community Structure. Duxbury and Haynie (2017) used network analysis to understand the structure of opioid distribution on the DNM Cryptomarket. As this DNM displayed buyer information, the authors were able to construct a network of opioid buyers and vendors from data collected on the 1 April 2016. This produced a bipartite graph with 706 buyers and 56 vendors as nodes connected by edges representative of 1132 purchases that took place between October 2015 and April 2016.

Network analysis showed that the network had extremely low network density, 0.2%, i.e. buyers tend not to branch out to new vendors when making purchases (Duxbury and Haynie (2017)). As a result, a small proportion of vendors are responsible for the majority of sales producing a high indegree centralisation. It was found that buyers purchase infrequently and so do not individually have a high influence over the network structure. No vendors were found to also be buyers (though it is unclear whether or not profile linking analysis was conducted) and almost a quarter of the vendors were entirely isolated from the graph, i.e. had made no purchases. As such, the study concludes that building a consumer base is difficult (Duxbury and Haynie (2017)).

36 unique communities were found with very little overlap, i.e. most buyers interact with a vendor who is relatively isolated from the rest of the network (Duxbury and Haynie (2017)). The vendors in larger communities tended to have high reputation scores and they were the only vendors to have buyers who made only one purchase. The nature of the structure implies that a buyer finding a trustworthy vendor does not necessarily encourage them to make more purchases (Duxbury and Haynie (2017)).

Vendor trustworthiness was measured using their cumulative reputation score, vendor affordability was measured using the average price of the products they have sold, and vendor diversity by the number of different opioid products they sell. These variables were used as predictors in Exponential Random Graph Model (ERGM) analysis to show that reputation is a stronger predictor for when a buyer will make a purchase than the relative price of the vendors products or which products they offer (Duxbury and Haynie (2017)).

This study was reliant on the available public information on buyers and, as such, cannot be replicated on other DNMs that do not present this information. As this information

is rarely made available on other DNMs this study will be difficult to replicate at all. Further, the authors chose to limit the study to opioid sales only and other products may exhibit different network characteristics.

In addition to looking at the DNMs themselves, other sites such as Google Trends and the DNM search engine Grams have been used to measure the popularity of different substances (Al-Imam and AbdulMajeed (2017a); Al-Imam (2017); Al-Imam and Abdul-Majeed (2017b)). These tools have been used to track interest in particular substances, such as Captagon, Octodrine, and NBOMe both temporally and geographically. The voting function built into Grams can also be used to evaluate the popularity of DNMs or vendors (Al-Imam and AbdulMajeed (2017a); Al-Imam (2017); Al-Imam and Abdul-Majeed (2017b)).

2.1.3 Law Enforcement and the Dark Net

There have been five large law enforcement efforts specifically targeting DNMs. The first was in October 2013 when the FBI successfully shut down *Silk Road* (Van Hout and Bingham (2014)) and arrested its creator Ross Ulbricht (U.S. Attoney's Office (2015)). The second took place in February 2014 when Dutch Police shut down the DNM Utopia in Operation Commodore arresting 5 people and seizing BTC900, or $610,900 (DeepDotWeb (2014f)). The third occurred a few months later in November 2014 when multiple law enforcement agencies conducted Operation Onymous and shut down several DNMs including *Silk Road 2.0*, *Pandora*, *Blue Sky*, and *Hydra* (Greenberg (2014b)). The admitted administrator of *Silk Road 2.0* Blake Benthall was arrested (U.S. Attorney's Office (2014b); The Ross Ulbricht Legal Defense Effort (2018)) alongside 16 others. The fourth effort, Operation Hyperion, took place in November 2016. This was also a multinational operation and involved the Five Eyes alliance (FVEY) comprised of Australia, Canada, New Zealand, the UK and the USA. Instead of shutting down DNMs, Operation Hyperion targeted individuals believed to have been active on DNMs and warned them to cease their activity (FBI (2016)). The most recent effort, Operation Bayonet, occurred in July 2017. The FBI and DEA seized the *Alphabay* which, at the time, was the largest and longest running DNM (Europol (2017)). In a seperate operation the Dutch police seized the DNM *Hansa*. They chose to keep the site running for information collection purposes and continued to do so until after the closure of *Alphabay* specifically in the hope they would increase the population of *Hansa* by attracting former *Alphabay* users (Europol (2017)).

The effect of these law enforcement efforts has been evaluated in a number of studies (Soska and Christin (2015); Lacson and Jones (2016); Aldridge and Décary-Hétu (2015);

Van Buskirk et al. (2015); Munksgaard et al. (2016); Décary-Hétu and Giommoni (2016); Dolliver (2015a,b); Van Buskirk et al. (2017),) the general consensus being that DNMs are resilient. For example, it has been shown that the DNM *Black Market Reloaded's* vendor population almost doubled and that *Sheep's* increased by 461% between the 7 November 2013 and the 3 October 2013 (Van Buskirk et al. (2014)) implying that vendors relocated after the fall of *Silk Road*, instead of leaving the Dark Net. Research conducted by (Soska and Christin (2015); Lacson and Jones (2016); Aldridge and Décary-Hétu (2015); Van Buskirk et al. (2015); Munksgaard et al. (2016); Décary-Hétu and Giommoni (2016)) implies that the arrest of Ross Ulbricht and the closing down of *Silk Road* did not meaningfully deter people from buying drugs online but, instead, may have facilitated the growth of other DNMs by removing their competition in a highly publicised event (Van Buskirk et al. (2014)).

It should be acknowledged that different results were found by Dolliver (2015a) when they scraped the site *Silk Road 2.0* on the 3 September 2014. They collected 1,834 listings, a much smaller number than the advertised 12,533 and found 199 unique vendors (Dolliver (2015a)). Dolliver (2015a) explains the discrepancy as the owners of the site potentially being deliberately misleading in order to make the market look bigger than it was or counting products that were not publicly available in their tally. It was also found that the drugs category was smaller than other categories such as eBooks differing significantly from the original *Silk Road*.

This study has been criticised as the findings differ substantially from other, similar studies indicating an issue with data collection and because of the premise that *Silk Road 2.0* should be directly compared to *Silk Road* to measure the impact of its take-down. *Silk Road 2.0* is potentially not the best DNM for comparison because it was not the strongest DNM in the ecosystem and it had recently suffered a large hack (Aldridge and Décary-Hétu (2015); Van Buskirk et al. (2015)). Indeed, a long term study whereby data was extracted manually, as opposed to with an automated scraper, collected 9,103 drug listings and 519 vendors on the 4 September 2014 (Van Buskirk et al. (2015)). A replication study was conducted on 9 partial crawls collected between the 4 August and 10 September 2014 taken from the Branwen dataset (Munksgaard et al. (2016)). 581 unique vendors and 12,259 unique items for sale were found, which describes a considerably different picture to that presented in (Dolliver (2015a)). A potential explanation for the discrepancy was presented: that the crawler used by Dolliver did not log error incidents, whereas the Branwen crawler did, therefore if an error had occurred during data collection, e.g. the crawler had gotten stuck on a vendor profile with a particularly large number of ebook listings, it may have crashed skewing the results in an undetectable manner (Munksgaard et al. (2016)).

A response to the criticisms of (Dolliver (2015a)) has been presented in (Dolliver (2015b)). Here, Dolliver (2015b) attempts to provide further evidence that *Silk Road 2.0* overestimated the number of listings advertised by including the results of a manual inspection of the site (Dolliver (2015b)). The response also calls into question the use of the Branwen dataset which was collected by an independent researcher and has not been peer-reviewed. Dolliver (2015b) claims that the "manually crawling approach" adopted by Van Buskirk et al. (2015) is also problematic as it will miss listings that are uploaded and removed during the time it takes to crawl the site. Finally, other, unpublished datasets cited in (Dolliver (2015b)) also point to *Silk Road 2.0* being especially volatile in nature before it was closed down and show that the number of listings varied by thousands from week to week. This volatility could potentially explain the contradicting depictions of *Silk Road 2.0* given by (Dolliver (2015a)) and (Munksgaard et al. (2016)) and allow for both studies to have accurately described the site.

However, empirical evidence in the form of police reports that describe the size of *Silk Road 2.0* after its closure shows that the data collected by Dolliver (2015a) is an underestimate. Indeed, new data presented in this body of work also demonstrates that *Silk Road 2.0* was bigger than Dolliver (2015a) claims, even at the beginning of its lifetime.

Analysis on forum discussions in the wake of *Silk Road's* closure by Lacson and Jones (2016) found that more comments discussed using *Silk Road 2.0* or other DNMs than stopping using DNMs altogether. They compared these discussions on the forums of several DNMs and found that the forums for *Agora* featured more comments of users who would not use *Silk Road 2.0* than the *Silk Road* forum (Lacson and Jones (2016)). They also found that the majority of users were waiting for more information before attributing a reason for the fall of *Silk Road* (Lacson and Jones (2016)). The overall sentiment of the messaging boards was positive or neutral.

When examining the forums of *Silk Road* and *Silk Road 2.0*, Horton-Eddison and Di Cristofaro (2017) focused specifically on the discussions around escrow technologies in the wake of ecosystem events such as the closure of *Silk Road* and *Silk Road 2.0's* hack. They found that the number of community discussions increased after the closure of *Silk Road* but the actual uptake of the technology began after a hack of *Silk Road 2.0* (Horton-Eddison and Di Cristofaro (2017)). As such, they concluded that the closure of *Silk Road* was a catalyst to the uptake of this technology.

Décary-Hétu and Giommoni (2016) evaluated the impact of Operation Onymous using a model designed to measure the efficacy of offline drug operations. They used data collected from the Gwern dataset on the top 5 most active markets prior to Operation Onymous. To measure any change in activity, the following factors were measured: the price per listing, collected weekly, the number of listings, the number and proportion of

active dealers as well as the number of new dealers who joined the markets studied that were not shut down in the operation and the number of displaced dealers, the overall drug consumption measured using the number of reviews (on a weekly basis), and the average concentration of dealers in the market (measured by the amount of feedback for each dealer divided by the total amount of feedback overall).

The data, as interpreted in this study, suggests that Operation Onymous had a positive but brief impact on DNM activity and a "chilling effect on the stable growth in the volume of sales" however, in the long term, it does not seem to have had a lasting effect. They conclude that crackdowns are not effective responses to online drug marketplaces and argue this is similar to offline drug markets (Décary-Hétu and Giommoni (2016)). It is worth noting, however, that a review of empirical studies by Pollack and Reuter, 2014, concludes that the risk of arrest, incarceration, or seizure in offline drug markets does not increase the prices of drugs (Pollack and Reuter (2014)).

Similarly, Van Buskirk et al. (2017) conclude that events, specifically Operation Onymous and the Evolution Exit Scam, may have immediate effects on the ecosystem population but do not reduce the growth rate (and therefore the population can recover). This conclusion was built on a longitudinal study examining the vendor population across 39 DNMs collected between October 2013 and November 2015 and used interrupted time series regression analysis to determine a significant change in population but not in rate of increase of population.

Work by Weisburd et al. (2006) has found evidence that crackdowns on offline drug markets can be effective, to the extent that they do not result in displacement to near by areas. This study also calls into question the variables used to measure the impact of Operation Onymous in (Décary-Hétu and Giommoni (2016)). Particularly as, instead of measuring displacement, the authors examine the number of new vendor accounts created on *Evolution* and *Agora* post Operation Onymous. They observe that, on these sites, the number of new vendor sign ups were increasing week on week before Operation Onymous but were decreasing between November 2014 and January 2015. Not only does this show a potential impact on the ecosystem from the operation but, further, these accounts are not explicitly linked to the accounts operational on the DNMs closed in Operation Onymous, so they are not necessarily displaced vendors.

Bhaskar et al. (2017) presented an overview of major DNMs from 2013 until 2016 and used the number of product listings as a measure for the impact of their closure. Whilst this measurement is not justified within their report, they demonstrate a similar pattern as with DNM populations of the number of listings and volume of product sales increases on DNMs that remain active after a Law Enforcement intervention or the closure of a DNM (Bhaskar et al. (2017)). They also evaluate the impact of the Evolution Exit

Scam on the overall ecosystem and find the market response is similar to that of a Law Enforcement intervention, i.e. minimal and short term.

Markopoulos et al. (2015) provide theoretical analysis as to how law enforcement could successfully manipulate the reviews of a Dark Net Market to reduce sales. They demonstrate, through optimisation problems in a simplified setting, the strategy that could disrupt a DNM by causing users to be unable to trust the review system and so leave (Markopoulos et al. (2015)). The analysis also shows that an operation with too few resources could have counterproductive results for example it could increase the number of sales (and therefore demand) on the market, reducing the overall quality of product and therefore put the health of customers at risk, or not have an impact but increase the number of sales overall when the programme ends (Markopoulos et al. (2015)). This paper does not provide analysis or argument explaining why reviews created by law enforcement will affect the behaviour of other consumers, nor does it provide real world evidence to confirm the analytical results.

Wadsworth et al. (2017) researched the impact of making NPS illegal in the UK. To do so, they measured the change in market share of NPS between the visible and hidden web (the Dark Net) before and after the introduction of the Psychoactive Substance Act. They found an increase in NPS listings on the hidden web and a decrease of availability in the visible web implying that, due to the change in the law, the sale of NPS had moved.

2.2 Related Contexts

Outside of DNMs, research has been conducted on the sale of "legal highs" online (Schmidt et al. (2011)), as well as other illicit items such as malware (Franklin et al. (2007); Motoyama et al. (2011); Van Wegberg et al. (2018)) and copyrighted material (Décary-Hétu and Giommoni (2016)). Whilst the facilitation of sales is different in these contexts compared to DNMs, the community structures may have similarities. Further, they present a similar problem to law enforcement and so research from this perspective, for instance that shows that removing phishing sites does have an impact on the industry even though some offenders set up replacement sites (Moore and Clayton (2007)), can inform studies on DNMs.

2.2.1 Legal Highs

Not all online drug sales take place on the Dark Web and research into sales on the Clear Web can supplement DNM knowledge, especially when more information is available

because of the lessened need for anonymity.

Schmidt et al. (2011) looked at legal highs available online in the UK. They searched, on Google and AltaVista, the term buy + "legal highs" + UK on the 7 April 2009. The first 100 pages and a random sample of 5% from the remaining pages were taken from each search, providing 115 pages which offered to sell legal highs in the UK, 39 of which were unique. Data was collected from these sites between April and June 2009. 1,308 products were found, they had an average price of £9.69 and took the form of pills, smoking material, and plant extract. The majority of the products were stimulants, sedatives, or hallucinogens with the top 5 products being Salvia divinorum, Kratom, Hawaiian Baby Woodrose Seeds, Fly Algaric, and Genie. Little safety, advisory, or ingredient information was available (Schmidt et al. (2011)).

2.2.2 Malware, Carding, and Cybercrime as a Service

Drugs are not the only illegal items available for purchase on DNMs, nor on the Dark Net as a whole. DNMs also facilitate the sale of malware, stolen information, and illegal services amongst other products. The literature on these cybercrimes, which considers both the DNM context, and specific cybercrime forums, can give insight on how cybercriminals consider anonymity and trust, as well as how they perceive law enforcement. Further, many of the methodologies and approaches in these areas can be applied directly to the DNM context.

Franklin et al. (2007) present an observational study about malware forums using data collected from January to August 2007 from International Relay Chat (IRC) (public group channels in this case used to discuss and sell malware). They collected 13 million messages of which 3,789 were selected uniformly at random and hand-coded into advertisements (by sellers and buyers). Additionally, syntactic analysis was conducted to identify regular expressions, e.g. credit card numbers and semantic analysis was conducted using a Support Vector Machine (SVM). These techniques were used to assess the proportion of the discussion which was actually advertisements and sales, and to distinguish between different types of product, such as between Social Security Numbers and credit card fraud.

Chu et al. (2010) explored 909 threads with 4,049 posts sampled from 6 online malware forums. They used qualitative analysis - Grounded Theory (Corbin and Strauss (1990)) and "normative orders" (Herbert (1996)) to understand and categorise their content. Posts were coded as advertisements or requests for particular products and these posts were subcategorised by the product being sold or sought. In addition the users posting advertisements were categorised, using Herbert's concept of "normative orders" (1996).

The orders determined to "structur[e] social interactions between buyers and sellers" were price, customer service and trust (Chu et al. (2010)). These were used to label sellers as "trustworthy, reputable, or cheat" using the comments made by them and those they interacted with.

Grounded theory was used to identify norms and values within the content, for example in relation to the establishment of trust on the forums. The data was all hand coded.

Motoyama et al. (2011) looked at 6 underground forums and their accompanying private messaging records. They analysed the degree distribution, social growth and user overlap then examined how these statistics correlated with sales and scams. The forums were modelled as networks with users as nodes and edges as representative of an interaction between two users either one user posting after another on a thread or sending a reciprocated private message. Users were linked from one forum to another by shared usernames and email addresses. Products and thread topics were sorted by searching for 500 manually created regular expressions. This study analysed the influence of social dynamics on trading by looking at how posting in a trading thread affected the number of interactions and ratings a user received. They then looked at the social dynamics around banning users and the status of the users being banned vs those doing the banning. The authors find that, where with Facebook 20% of a user's peers a responsible for 70% of their interactions, on these forums 70% of a user's peers are responsible for 70% of their interactions (Motoyama et al. (2011)).

Soudijn and Zegers (2012) conducted text-based analysis on 15,000 posts from a carding forum (a forum for the buying and selling of stolen credit card information). This analysis was used to construct crime scripts of the carding process and identify interventions through the application of Situational Crime Prevention (Soudijn and Zegers (2012)). They found the elements of the crime (leaving traces whilst wiring money) which were of greatest concern to the offenders themselves (Soudijn and Zegers (2012)).

Décary-Hétu and Giommoni (2016) looked at the Warez community, a group that shares copyrighted content online illegally. They looked at 5 large scale law enforcement operations that resulted in 175 convictions and used a web crawler to measure the number of uploads still available after each. Their analysis showed no real change in activity after the operations, leading them to conclude that large scale law enforcement efforts of this nature are not worthwhile. This was to be predicted because the ecosystem is very large and volatile, with a high turnover of players, these players are in direct competition and so benefit from law enforcement removing their competitors, and, they are insulated by anonymity so most participants are not affected unless they are the small proportion directly affected (Décary-Hétu and Giommoni (2016)).

Van Wegberg et al. (2018) use scraped datasets of DNMs to measure the cybercrime-as-a-service products also available on these sites. The authors measured the different forms for business-to-business products (e.g. botnets, e-mail accounts, malware, and cash-out services) and business-to-customer products (e.g. fake documents, pirated media, and how-to guides) across 8 DNMs (*Agora*, *Alphabay*, *Black Market Reloaded*, *Evolution*, *Hydra*, *Pandora*, *Silk Road*, and *Silk Road 2.0*). It was found that business-to-business cyber crime generate approximately $8 million and business-to-customer cybercrime approximately $7 million between 2011 and 2017 (Van Wegberg et al. (2018)). Both figures are likely to be underestimates and are limited by the data collections technique (Van Wegberg et al. (2018); Christin (2013)).

The data was scraped by Christin (2013) and reviews were used to approximate sales. The product categories were determined by identifying the key stages of a cybercrime that could be outsourced and each product was categorised using a Linear Support Vector Machine (SVM) classifier which was trained on 1,500 hand coded products, though examples from underrepresented categories had to be added to the ground truth to avoid biasing the classifer. The average precision of the classifier was 0.78 and the average recall was 0.76 (Van Wegberg et al. (2018)).

It was found that the number of listings, the number of pieces of feedback, and the total monthly revenue rose steadily between 2011 and 2014 before increasing steeply and then entering a period of volatility with several sharp decreases (attributed to Operation Onymous, the Evolution Exit Scam and the closure of *Agora*, though not statistically) before returning to a steady rate of increase more rapid than at the beginning of the measurement period (Van Wegberg et al. (2018)). The vast majority of revenue was associated with cash-out services as were the majority of vendors (Van Wegberg et al. (2018)).

Latent Dirichlet Allocation (LDA) was used to extract key products for each category. This approach was taken so that key products could be determined from the review data, as opposed to simply looking for products that generated the most revenue - a technique which is biased towards expensive products. For cash-out services, the top three clusters were fake credit card details, stolen credit card details, and guides for recruiting money mules (Van Wegberg et al. (2018)).

2.2.3 Hackers

The hacking community bears some resemblance to DNM users as both groups are centred around an illegal activity that requires at least some technical knowledge in OPSEC

and could be motivated either by profit or a political belief. Research on the motivations and methods of hackers, therefore, can provide a foundation for understanding DNM users especially where studies in this discipline can be replicated. Some methods of conducting sales are also utilised by both communities and therefore methodologies developed to understand hacker forums can also be applied to the DNM context.

Caines et al. (2018) developed tools for the automatic classification of posts on hacking forums in the CrimeBB corpus (a large collection of posts from English and Russian language forums including HackForums). Three annotators read over 2000 posts from randomly chosen messaging boards and coded them according to type, addressee and author intent.

These hand coded posts were used as a training set for multiple logical and statistical models. They find that a hybrid logical-statistical classifier best labels posts by type and author intent but a statistical classifier best labels the addressee of the posts. The models were demonstrated to be efficient and scalable for the remainder of the dataset (Caines et al. (2018)).

Whilst there was some ambiguity in the labelling of some posts (for example, some posts did not appear to address a specific person or might be responding to a comment not directly preceding it), there was strong inter-annotator agreement which likely aided the accuracy of the models built.

Pastrana et al. (2018) argue that few users of underground forums participate in illegal activity. They conducted analysis on *Hackforums*, the largest forum in the CrimeBB dataset which contains 30 million posts from 572 thousand users, to identify and predict those key actors likely to engage in illegal activity.

Before analysing the data, a ground truth of key actors was assembled. They looked for actors who were of interest to law enforcement by using media sources to see who had been arrested, receiving intel from a private security company, finding these actors' closest neighbours, identifying the actors advertising top Remote Access Trojans (RATs). In total they determined 130 actors of interest, 113 of which could be found in the forum.

Using k-means clustering, where $k = 5$, they extracted 44 features for each actor. These features relate to their forum activity (e.g. number of days between their first and last post, etc.), network centrality measures (e.g. out and in-degree, eigenvector centrality, etc.), and reputation measures (e.g. overall reputation score calculated using the reputation system on the forum, etc.).

Logistic regression was used to build a prediction model (informed by the extracted features) for determining who in the forum will become a key actor during its lifetime.

The predictor had a false positive rate of 0% and accurately identified 12 out of the 108 (5 were removed because of "undue influence on the model") key actors. The authors argue that the model is an effective tool, despite its low success rate, because it identified those 12 actors from a set of nearly a quarter of a million and because forum activity is unlikely to be the only source of information employed when identifying key actors (Pastrana et al. (2018)).

The prediction analysis was then repeated using the topics discussed by the key actors. A combination of logistic regression, social network analysis, and clustering approaches were used to determine other actors of interest within the dataset. In total 80 actors (from a possible 285) were identified as having similar profiles, interests and social behaviours as the 113 key actors identified at the beginning of the study.

Young et al. (2007) conducted a survey at DefCon ("the largest annual computer hacking convention" Young et al. (2007)) to determine the attitudes of hackers towards law enforcement and illegal hacking. The survey, which was completed by 127 participants across the 3 day conference, was distributed to self identified black hat hackers who claimed to have broken the law using hacking techniques in the 12 months preceding the conference, students (who had not broken the law) and other attendees. The responses of black hat hackers were then compared to the responses of the other two responder groups.

The survey was designed to measure the moral disengagement (how illegal hacking is morally justified by the respondent), informal sanctions (how the respondent feels their community would treat them if it was discovered they participated in illegal hacking), punishment severity, punishment certainty, and utility (what the respondent felt they gained from hacking). The study concluded that the hackers were able to justify their behaviour, often by blaming their victims for their actions, and considered hacking to have a high utility value and few repercussions (Young et al. (2007)). Indeed, the hacker respondents believed they would not lose support from friends or family if they were found to be hacking and, despite recognising potentially high legal repercussions, felt that they had a low chance of being caught (Young et al. (2007)). Further, hackers perceived the likelihood they would be caught to be statistically significantly lower than the perceived likelihood for the student and other attendee populations (Young et al. (2007)).

The discussion of results in this study was also informed by the model presented in (Kshetri (2006)) which describes "the viscous cycle of cybercrimes". Kshetri (2006) describes how the characterisation of law enforcement as ill-equipped and out matched by cybercriminals both informs and is perpetuated by the decision not to report attacks due to a lack of faith in law enforcement. Further, these both in turn increase cybercriminals'

confidence and success rates which feeds the perception, and reality, of law enforcement being unable to tackle the problem (Kshetri (2006)).

Afroz et al. (2013) consider the sustainability, as opposed to profitability, of hacking forums. They apply Ostrom's economic framework of commons governance to 5 hacking forums (AntiChat, BadHackerZ, BlackhatWorld, Carders and L33tCrew) which have varying success - success of the forums is measured by the presence of small world characteristics in the network as these networks have been established to be more commercially successful, economically efficient, and creative (Afroz et al. (2013)). A network is considered small world if it is highly clustered and has a small path length and this was the case for 4 out of the 5 observed forums.

Ostrom's economic framework of commons governance argues that the sustainability of a resource within a community is dependent on the community meeting the following 5 criteria: "1) low cost of monitoring, 2) moderate rates of change of the resource, 3) frequent communication between resource members, 4) low costs of enforcement, and 5) exclusion" (Afroz et al. (2013) pg.3). By examining the network structures of the networks created from public and private messages between users of each forum and by understanding the rules and enforcement mechanisms of each forum, Afroz et al. (2013) argue that the 5 criteria outlined above "correlate with successful underground forums online" (Afroz et al. (2013) pg.8). Though, it should be noted that the sample size of this study was small and conclusions were not demonstrated statistically.

2.3 Resilience Theory

Resilience Theory is, broadly, a framework for understanding how systems or individuals respond to negative events and measuring how capable they are of surviving those events. Different disciplines have different specific definitions and models of resilience theory developed to encapsulate the particular negative events relevant to them. These definitions and tools can be applied to the problem of DNMs as they may help to identify the characteristics that enable the ecosystem to survive law enforcement interventions and other events. Pin pointing such characteristics could inform law enforcement strategy such that future efforts attempt to weaken the resilience of the ecosystem before, or as well as, shutting down DNMs and arresting individuals. This section describes the different concepts of resilience within the disciplines of Psychology, Ecology, Sociology, and Engineering as well as detailing some specific, relevant models that informed the methodology and research direction of this thesis.

2.3.1 Psychology

Within psychology, resilience is often focused on individuals, families, or, more recently, communities and is defined to be "positive adaptation despite experiences of significant adversity or trauma" (Luthar (2006) pg.742). The methodologies for identifying resilient attributes of individuals in high-stress environments can be applied to DNMs in order to understand which characteristics make a vendor more likely to continue trading even after a DNM closure or law enforcement intervention. Concepts and models within this area may also be useful for defining resilience in vendors, beyond simply measuring if a vendor is trading or not.

An extension of the concept of resilience is that of thriving, when an individual is not simply resilient to an adverse event but able to improve themselves through the process of experiencing the event (Van Breda et al. (2001); Carver (1998)). If the individual acquires new skills or knowledge that leave them better equipped to deal with new events they have thrived in the event. Applying this concept to the study of DNMs may allow law enforcement to identify adaptations in behaviour which make future interventions less impactful and adapt their approach where current interventions enable or encourage these adaptations.

The psychology literature which focuses on community resilience describes how communities are able to survive or rebuild after disastrous or damaging events. If applied to the DNM ecosystem, the methodologies developed to achieve this goal may be used to understand if the ecosystem has survived large scale law enforcement interventions. Further, were DNM communities to behave like the physical communities that are studied within this field, their conclusions may help to explain why the ecosystem persists.

2.3.1.1 Resilience in Individuals

As opposed to observing resilience directly, it can instead be inferred based on the presence of significant adversity (or risk) and positive adaptation (or competence) (Luthar (2006)). Thus, measuring resilience, or the propensity to be resilient, can involve identifying the adversity an individual is exposed to and measuring the evidence of competence. Adversity is measured using risk factors, which are considered to be the aspects of an individual or their environment that make them vulnerable to maladjustment.

Risk factors are often defined operationally and can be thought of as predictors of bad future outcomes (Masten (2001)). Within the context of child psychology exposure to violence or maternal postnatal depression would be considered risk factors (Luthar (2006)). Risk factors are often defined as statistical probabilities and can be considered

in isolation or, preferably, as cumulative factors that intersect and produce compounded results (Luthar (2006)). Whilst examining multiple risk factors in combination may provide a more realistic picture, testing risk factors in isolation can help to identify specific points for intervention and it is therefore important to consider both approaches (Luthar (2006)).

Within the literature, different ways of combining multiple risk factors have been presented. Composites can be formed by summing standardised numerical values representative of the amount of risk an individual faces in each dimension, where each dimension is a different risk factor (Masten et al. (1990)). Alternatively, risk factors could be presented as binary - either the individual is at risk or is not and a metric is used to translate continuous risk factors, such as intelligence scores, into binary values. In (Gutman et al. (2003)) an individual was determined to have a risk factor if they were in the top quartile of the population. The total risk of an individual can then be presented as the sum of the assigned values.

Positive adaptation is measured by observing competence considered to be unexpectedly better than the present risk factors would leave one to expect (Luthar (2006)). In order to define "unexpectedly better than", the criteria set must be developmentally appropriate and relevant to the risk factors being measured, for example, if a subject carrying many risk factors for antisocial behaviour displays socially conforming behaviour, this may be an indication of a positive adaptation (Luthar (2006); Seidman and Pedersen (2003)). An alternative approach is to define competence based on expected behaviour, for instance, setting out cultural age expectations and measuring whether a subject has met them (Masten (2001)). As with risk factors, competence can involve multiple dimensions and so should be measured with that in mind. It is important to consider when a subject is displaying positive adaptation against one risk factor but failing to do so against others (Luthar (2006)). Similarly to the methods described above, competence can be measured individually or producing a composite by standardising the measurements of different factors (Bolger and Patterson (2003)).

Positive adaptation can also be referred to as a resilience quality (Richardson (2002)) and this is often identified by observing groups in high risk environments and identifying the qualities unique to, and shared by, the participants who did well despite their risk factors or by determining qualities necessary to thrive and testing to see if they were found in the subpopulation of participants who thrived in an high risk environment. As DNM users are potentially always at risk (for example of being arrested) this approach may be applicable to studies of the ecosystem. For instance, by subcategorising users into groups that share certain qualities hypothesised to be positive adaptations.

There are two different ways of measuring risk factors and competence (Luthar (2006)). The first is variable driven - risk factors and positive adaptations that protect against those factors are hypothesised and statistical analysis is performed to evaluate whether or not the risk factor measurements correlate with the competence measurements as predicted (Garmezy (1987)). The second approach is to compare different groups of subjects in the 2-dimensional space of competence and risk, e.g. subjects who display high competence and have high risk, vs, subjects who display low competence and have high risk, etc. (Luthar (2006)). This then allows for the identification of positive adaptation which distinguishes the different groups. There are two dominant methods for sorting subjects into groups. The first produces composite values of risk and competence as outlined above and then sorts subjects into predefined groups based on their score (Luthar (2006)). The second involves creating cut off points (normally defined by mean values or clinical standards) whereby subjects that are above the cut off are sorted into one group and vice versa (Luthar (2006)). This person-based analysis not only allows you to compare resilient and non-resilient groups but also other dimensions, such as specific competence based criteria.

The limitations surrounding studying the DNM ecosystem may make developing scales of risk difficult. For example, the fact that most users interact with the ecosystem anonymously and take steps to hide their operations mean that understanding a user's resilience with a high level of complexity is difficult. From a practical perspective, simpler binary measures (e.g. does this user have this characteristic?) are likely more appropriate for understanding resilience.

A resiliency model is presented by Richardson (2002) to describe how individuals can react to events, it proposes that there are four outcomes: resilient reintegration, reintegration back to homeostasis (a relatively stable equilibrium), reintegration with loss; and, dysfunctional reintegration. This means that, from the disruption, they can improve on their current state beyond what it was before the disruption (i.e. thrive); return back to their state before the disruption; return to their previous state but with less motivation to reintegrate in the future; or, adopt new, destructive behaviour. Resilient qualities, or positive adaptations, factor into the model by determining whether or not a disruption will occur and, if so, on what scale.

Similarly, in the DNM ecosystem, there are multiple ways of engaging with the system. Users can actively engage (by buying or selling) or passively engage (by owning but not using an account). If these levels are evaluated differently this might reveal more nuanced ways that users are affected by ecosystem interventions.

2.3.1.2 Community Resilience

The understanding of an individual's resilience being dependent on their community has led to the development of community resilience in addition to models of resilience for individuals. Magis (2010) defines community resilience to be the "existence, development, and engagement of community resources by community members to thrive in an environment characterised by change, uncertainty, unpredictability, and surprise" (Magis (2010) pg.402).

Norris et al. (2008) present a model for community resilience after a disaster. The model presents three outcomes, similar to the model for individual resilience found by Richardson (2002). They are *resistance*, where the negative impacts of the event are avoided; *resilience*, where the community adapts to the circumstances of a crisis and continues to function; and, *vulnerability*, where the community fails to adapt, or adapt fully, and becomes dysfunctional (Norris et al. (2008)). It is stressed that resistance is the least likely occurrence (Norris et al. (2008)). The determinants of the outcome are the severity, duration period, and surprise factor of the crisis and the robustness, redundancy, and rapidity of deployment of the community's resources. These can be considered as the risk factors and competencies of a community. These factors may also be evaluated when analysing the ecosystem's response to large scale events.

When seeking to measure the adaptation of a community, Norris et al. (2008) consider a community more than the sum of its parts and so measuring the community is more than summing measurements of the individuals within the community. They present a framework whereby the resilience of a community is measured by its adaptive capabilities, i.e. how robust, redundant and rapidly deployed its resources are.

The DNM ecosystem may be considered from the perspective of the individuals who use it or as one community. These two perspectives may be linked but should be explored both separately and together to understand if the approach of Norris et al. (2008) is appropriate.

2.3.2 Ecology and Sociology

The sustainability and resilience of systems has been a primary focus of Ecology for many decades and, as such, much work has been done to model complex ecosystems which are affected by varied, and often random, factors. If the DNM ecosystem and its users behave like ecological systems, for example like coral reefs and their surrounding sea life, then the factors that make those ecosystems resilient or vulnerable to attack may have equivalents within the DNM ecosystem. Even if this is not the case, the

literature on ecological and sociological resilience provides methodologies for defining and understanding resilience that may help to develop a framework of resilience in DNMs.

Holling (1973) presented a framework for modelling the interactions of species in such systems and proposed measuring resilience as the likelihood of certain species going extinct. This measure comes from trying to determine how able the system is to return to its former state before a disturbance event. A concept which has been built upon to produce a definition applicable to social-ecological systems – resilience is the capacity of a system to retain its identity or primary function despite change or disturbance (Walker et al. (2004); Cumming et al. (2005)).

This definition allows for extinction and other major changes as long as the system's primary purpose is preserved. As such, when determining whether or not a system is resilient, the system's primary purpose and deviations from that purpose must be defined as well as disturbances to the system (Ludwig et al. (1997)). This approach can be applied to the DNM ecosystem by setting the types of law enforcement disruption and the objectives of the system from a user perspective.

Cumming et al. (2005) present a 5-step research design to measure the resilience of a system and identify ways to improve it and protect the system from adverse change. First, the identity of the system must be defined, then alternate future systems must be described, then the change trajectories determined, this allows for probabilities to be assigned to future scenarios and, finally, for interventions to be identified (Cumming et al. (2005)). If applied to the setting of DNMs, this system could be used instead to identify ways of weakening the resilience of the overall ecosystem, thus making law enforcement intervention more effective.

The foundation of the system's identity should come from the inhabitants of the system, for example through consultation or workshops. From their descriptions of system, the basic infrastructure such as location and temporal space can be defined as well as the essential system attributes. In addition, any variables that are likely to change in response to external or internal drivers can be listed and filtered for relevance to the research question. Each of these can be combined to produce a model of the system (Cumming et al. (2005)).

The second step is to define future systems. These systems can either retain the identity of the original system or be entirely different, or something in between. If no future alternatives can be imagined, this is an indication that the system is not resilient to change (Cumming et al. (2005)). In conjunction with this step is the third step, which involves explaining how the existing system would move to the possible, future systems.

The purpose of describing the trajectories of change which move the system from one description to another, is to determine how the key aspects of the system's identity affect its resilience (Cumming et al. (2005)).

The fourth step involves using quantitative analysis to assign probabilities to each of the future scenarios. It determines which are likely to occur and which aspects of the system's identity will be retained regardless of which scenario is moved to over time. From this, the resilience of the system can be determined as the likely change in identity can be described (Cumming et al. (2005)). Finally, the fifth step involves identifying and prioritising interventions that reduce the likelihood or guarantee specific types of change.

The use of these definitions and models of resilience have helped to identify a number of resilience factors in different social and ecological contexts. For instance, Côté and Darling (2010) examined the role of local stressors in determining the resilience of coral to climate change. They showed that local stressors, such as overfishing, were able to build resistence behaviours in coral making it more resilient to climate change and that an absence of local stressors after climate change damage enabled a greater capacity for recovery (Côté and Darling (2010)). Elmqvist et al. (2003) also examined the case study of coral reefs and found that a diversity of species within the reef can aid in resilience. Research on species dominance in Western Polynesia has shown that species which rely on animals to disperse their seeds were able to maintain species dominance after events that dramatically reduced the diversity with the group of pollinators because they were pollinated by multiple species (Elmqvist et al. (2003)). An examination of the role of plant species in ecosystem resilience conducted by Walker et al. (1999) found that the long tail distribution of plant species biomass was a resilience feature. It was observed that a few species dominate the ecosystem whereas many other species are only represented in small samples (Walker et al. (1999)). Further, the dominant species are diverse in their characteristics whereas the less represented species are likely to share characteristics with at least one dominant species (Walker et al. (1999)). This means that the dominant species are not fighting for resources and if one is unable to survive an environmental stressor or disturbance, it can be replaced by less dominant species (Walker et al. (1999)).

Ayling (2009) has applied the resilience theories found in ecology and sociology to crime networks, specifically gangs, to identify characteristics that make organisations resilient to law enforcement efforts. Thick crime habitats, community support, and interpenetration were identified as factors that bolster gang resilience, i.e. the availability of environments rich in crime targets and other offenders, the protection of a non-criminal community, and owning a large network of legitimate and illicit businesses all make it

difficult to break down criminal gangs (Ayling (2009)). An examination of the organisational structure, and organisational resilience literature, also identified specific organisational elements that aid resilience. They were *semi-structures, empowered members, and shared vision*, these characteristics allow for the gang to adapt quickly to changes in the environment as the gang is not structured hierarchically, which means that individual members have the ability to make decisions but are likely to make those decisions prioritising the same goals as other members; *redundancy*, this ensures individuals in the gang are easily replaceable if they are arrested; *hubs, weak links, and loose coupling* this is a structure that allows for sections (or hubs) of the gang to be destroyed by law enforcement without it damaging the rest of the gang thus minimising the effect of law enforcement efforts; *gang history and strong ties* the culture of the gang in terms of how it was created and the identity it provides to its members may ensure greater trust and a heightened capacity to follow orders; *secrecy and compartmentalisation* similarly to the hub based structure, the compartmentalisation of secrets can help to minimise the damage done when members of the gang are compromised; *bricolage* this is the ability to make use of available resources when it is not immediately apparent how they may be beneficial and is facilitated by the semi-structure and hub based structure but can be restricted by secrecy and compartmentalisation; *organisational learning* this is the capacity for a gang to learn and adapt as well as disseminate this knowledge – effective dissemination of knowledge can help to improve resilience where that knowledge is based around, for example, new evasion techniques (Ayling (2009)).

Many of these qualities have also been identified as improving resilience in cities vulnerable to hazards like natural disasters or terrorism. Godschalk (2003) collated these qualities to produce the following list:

- *redundant* there are several components that serve the same function;
- *diverse* there different components that protect against different threats; *efficient* the energy supplied is greater than the energy used;
- *autonomous* different components can operate independently when necessary;
- *strong* the city actually has the capacity to resist hazards;
- *interdependent* the components of the city are able to support each other;
- *adaptable* the city can gain resilience qualities during and after surviving a hazard;
- *collaborative* the city provides multiple opportunities and incentives for stakeholder investment (Godschalk (2003)).

Given how this approach has been translated to different contexts it provides a useful framework for identifying the key resilience factors within the DNM ecosystem.

2.3.3 Engineering

Engineering Resilience contrasts to resilience in Ecology both in terms of its focus and measurement – for Engineering Resilience the focus is on the efficiency with which a system returns to equilibrium after an event, as measured by the time taken or rate at which it returns whereas, for the resilience of a system in Ecology, the focus is on the existence of any function and resilience is measured by the amount of disturbance that can be experienced until function is lost (Reggiani et al. (2002)). Given this, it is harder for notions of resilience in Engineering to capture a constantly evolving environment (Reggiani et al. (2002)), such as the DNM ecosystem.

Resilience in Engineering has been defined as the length of time taken for a system to recover and how efficient that process is (Folke (2006)). It focuses on the recovery of the original system and how constant the state of the system is. For ecological systems that are constantly changing, and not always in response to harmful events, this focus is not necessarily helpful (Folke (2006)). We may consider the DNM ecosystem to be more similar to an ecological system than a machine and so the concepts of resilience found within Engineering are less applicable. However, there are still some useful frameworks, particularly within the discipline of Networks where resilience is defined as the ability to maintain service, or the continued provision of access to information, during and despite adverse events (Madni and Jackson (2009)).

Madni and Jackson (2009) provide a framework that presents resilience as a multifaceted capability which includes the following: *avoid*, the ability to predict disturbances and make necessary arrangements to ensure the system is unaffected; *withstand*, the ability to remain robust against disturbances that cannot be avoided; *adapt to*, this is the ability to reconfigure the structure due to a disturbance; and, *recover from*, the ability to restore the system to its pre-disturbance state.

Utilising this framework to assess resilience involves considering all of the events that pose a risk to the system and determining what system components are able to avoid, withstand, adapt to, or recover from the event. To do this, a thorough understanding of the systems attributes on multiple levels (from the actual components to organisational infrastructure) must be articulated and a variety of different methods of risk management and assessment must be employed. Madni and Jackson (2009) propose measuring resilience using the metrics of the time and/or cost taken to restore the system operations, the degree to which the pre-event system can be restored, the severity of any potential disruption which is circumvented, and/or, the successful adaptations gained in response to a disruptive event.

A primary component which makes engineering systems resilient is the accurate prediction of all potential threats and constant adaptation in order to prepare for those threats (Hollnagel et al. (2007)).

More recently, resilience in engineering has been defined to account for changes in a system. For example, Hollnagel et al. (2007) considered resilience as the "ability of a system or an organisation to react to and recover from disturbances at an early stage, with minimal effect on the dynamic stability" (Hollnagel et al. (2007), (pg. 16)). This is because systems change over time (are dynamic) and can withstand small disturbances that impact on the stability temporarily, without becoming entirely unstable. Another way of considering the temporal aspect of resilience is in the definition "resilience is the ability to prevent something bad from happening, or the ability to prevent something bad from becoming worse, or the ability to recover from something bad once it has happened" (Hollnagel et al. (2007), (pg.59)) which recognises that a system can be negatively impacted by an event but still be resilient if those negative impacts were not the most negative they could be, or are recovered from.

Similar approaches have already been taken in the literature on the DNM ecosystem when the time taken for new DNMs to emerge after interventions is measured.

Risk assessments can be used to make predictions on when a system loses its stability. This can involve looking at the events and functions in the system, not just the components.

In large, complex systems, resilience can be created through the interactions of components that cannot be provided by the components individually (Hollnagel et al. (2007)). This complexity can make a system resilient but also make the resilience of the system difficult to measure. A suggestion for how to determine the resilience of the system, is to look at the events and functions of the system, as opposed to the individual components. Conducting a risk assessment can lead to predictions on how functionality may be lost due to specific events and therefore give an indication of a system's resilience. Alternatively, Measures of resilience can be found by observing the system as it is stressed, e.g. by "abstract[ing] general patterns from specific cases of challenge and response" (Hollnagel et al. (2007)).

A framework that seeks to identify where the differing functions in a complex and/or dynamic system can interact to make the system less resilient is the Functional Resonance Accident Model (FRAM) (Hollnagel and Goteman (2004)). In this model, the functions of a system are described in relation to each other using a specific framework that highlights the relationships between functions. The goal of the framework is to identify relationships between functions that compound negatively when functions do

not occur as intended. This is done by visualising the system through its functions and their relationships.

The relationships are considered in terms of the inputs, outputs, resources, controls (or constraints), preconditions (conditions that must be fulfilled before a function is carried out), or time (how long a function takes or when it must take place in the system) of each function. For instance, if the output of one function is a resource for another function, they would be connected from output to resource. This means that, if the first function fails, the second may also fail, as may all other functions that have a relationship with it.

Given the complexity and interconnected nature of the DNM ecosystem, this approach of breaking it down into functions which are expressed in relation to inputs and outputs of the system may lead to a greater understanding of its resilience. It could also make it easier to identify variables that can be measured as indicators that an event has harmed or bolstered the ecosystem.

Chapter 3

Research Question and Hypotheses

The Research Question in this dissertation is as follows: *Is the DNM ecosystem resilient to law enforcement interventions?*

In this section, a definition of resilience will be justified, then the events that the ecosystem must be resilient to are outlined and, finally, hypotheses about the resilience of the ecosystem given the occurrence of such events will be presented.

3.1 Resilience Definition

In order to answer the research question, a definition of resilience must first be presented and the definition found in (Cumming et al. (2005); Walker et al. (2004)) i.e. resilience as a measure for how able the ecosystem is to retain its primary function despite change or disturbance, is used. The primary function of the ecosystem, according to its users, is to buy and sell drugs with greater convenience and safety than found offline (Munksgaard and Demant (2016); Van Hout and Bingham (2013b,a, 2014); Van Buskirk et al. (2016b); Barratt et al. (2014)). Convenience includes a wider availability of products to choose from, a reliable and understandable method of evaluating vendor and customer competencies, privacy technologies that can be learnt (potentially with some effort, e.g. PGP encryption), a large potential customer base, and the potential for a high mark-up for vendors. Safety includes a higher quality of drug and a low risk of being caught and arrested.

This gives multiple metrics of resilience when examining the ecosystem, each of which has been measured in other studies that evaluated the impact of law enforcement on the ecosystem:

- the size of the population (Van Buskirk et al. (2014); Soska and Christin (2015); Lacson and Jones (2016); Aldridge and Décary-Hétu (2015); Van Buskirk et al. (2015); Munksgaard et al. (2016); Décary-Hétu and Giommoni (2016); Dolliver (2015a));
- the quantity and variety of available products (Dolliver (2015a));
- sales revenue (Soska and Christin (2015));
- product prices (Décary-Hétu and Giommoni (2016));
- the adoption of privacy-enhancing technologies (Soska and Christin (2015); Broséus et al. (2016)).

This definition of resilience has been chosen because, even if the ecosystem contains entirely different DNMs and a whole new population after an event, if it is still functioning then it is still of interest to law enforcement. Whereas, if an event occurred that left users operating and DNMs online but the ecosystem was no longer used to trade illegal goods, law enforcement may no longer be interested. Similarly, users of the system are not necessarily willing to trade in all circumstances and so the ecosystem may be considered not resilient to events that decrease the convenience and safety of using it.

A dimension of the ecosystem resilience is the resilience of its users. This is because, even though the ecosystem may retain its resilience with new users, the resilience of the users can contribute to ecosystem resilience. This may be because the users themselves are resilient or because their collective behaviour creates a resilient environment.

The resilience of the user is defined as in (Luthar (2006)), i.e. in terms of the competence an individual user displays in the face of their risk factors. The risk factors of a user are the characteristics that make an individual vulnerable to maladjustment. In this instance, an intuitive definition of a maladjusted user would be one who is unable to trade successfully (as either a vendor or buyer), because the purpose of the DNMs is to facilitate the trade of products (Munksgaard and Demant (2016)).

However, as discussed, interviews with users imply that the attraction of the ecosystem is not merely the ability to trade but also the convenience of the trade, the quality of the products, the amount of demand, and the profit margins (Van Hout and Bingham (2014)). A user who is unable to obtain these properties could also be considered maladjusted. This gives several measures for maladjustment:

- whether the user is trading or not;
- the user's description of the quality of the trade;

- the number of trading partners available to the user;
- the frequency of trade;
- the cost of trade.

None of these are objective measures but, instead, can be measured comparatively, i.e. if a user loses trading partners or their frequency of trades decreases, this can be an indicator of maladjustment.

Competence is "effective functioning in important environments" (Masten et al. (1992), pg.239) and can be measured using qualities found in users who are functioning effectively (Masten et al. (1992)). If a maladjusted user is defined as one with no or diminished trading capacity then an adjusted, or competent, user is one with full trading capacity. Given this definition, the measures used to assess maladjustment can also be used to assess competence, i.e. users that maintain their

- trading status;
- quality of trade;
- number of trading partners;
- frequency of trade;
- profits from trade

can be considered to be competent.

Therefore, the resilience of users in the system is a measure of their competence (their ability to maintain their interaction with the ecosystem) given the risk that they could become maladjusted. Whilst all users are at risk of becoming maladjusted they are at a greater risk if they are affected by an adverse event or present whilst an adverse event takes place on the ecosystem. These such events and the way they affect the ecosystem and put users at risk of maladjustment are described in the next section.

3.2 Adverse Events

The ecosystem must be resilient to many different types of adverse events that may affect its ability to function. Some of these events are caused by law enforcement with the intention of damaging the ecosystem and are as follows:

- *DNM Shut Down*: this is any event in which a, or multiple, law enforcement group(s) is able to shut down at least one DNM and seize the server and any cryptocurrencies belonging to the DNM, for example the closure of Silk Road

(Christin (2013)), Operation Commodore (DeepDotWeb (2014f)), Operation Onymous (Greenberg (2014b)), and Operation Bayonet (Europol (2017)). These events are identified by law enforcement claiming responsibility for the DNM closure.

- *User Warning*: this is an event such as Operation Hyperion (FBI (2016)) in which DNM users are approached by law enforcement and warned of the potential for arrest but are not actually arrested.
- *Arrest*: alongside law enforcement efforts targeting whole DNMs, some operations have focused on arresting individuals or groups of buyers and vendors. These have occurred on an ad-hoc basis and in large coordinated efforts, such as Operation Hyperion. To identify arrests, arrest records have been used as well as news stories like (DeepDotWeb (2015e)) and work by community members and researchers to link arrest records that do not state the username of the defendant to active vendors (Branwen (2017)).
- *Parcel Seizure*: not all of the parcels shipped make it to their destination (Reddit User (2016)) as law enforcement attempt to identify and intercept parcels containing illicit items. Seized shipments are identified through reviews and comments and so are based on DNM user speculation, as opposed to definitive proof.

The ecosystem may also be damaged by its own users who could have motivations that conflict with the goals of the ecosystem. These are:

- *Closure*: DNMs are not always profitable or otherwise worthwhile for the admin team to run. This fact, and, other priorities they might have, may cause an admin team to declare they will close the site and then do so.
- *Exit Scam*: another reason a DNM might close is that its admin team chose to shut it down without warning and with the intention of stealing any coin held in escrow. Whilst it cannot always be known for certain if a DNM is closed in an exit scam or not, exit scams usually occur without warning to prevent users removing their coin, as such they are defined as the DNM closures for which the admin team did not issue a formal warning. A notable example is the clousre of Evolution in March, 2014 (DeepDotWeb (2015d)).
- *Hack*: rival DNM admin, vigilantes trying to expose bad practice, financially motivated hackers and even greedy admin have also hacked into DNMs and stolen coin held in escrow (Reddit User (2014e); DeepDotWeb (2014d); Reddit User (2014g); DeepDotWeb (2015c, 2014c).) The hack of Silk Road 2.0 in February, 2014 resulted in a loss of $27 million (DeepDotWeb (2014d)). Hacks are identified either through the hackers themselves announcing their actions or site admin claiming to have been hacked.
- *Scam*: vendors may consider it more profitable to accept payment from buyers without shipping any product, similarly buyers may claim to have not received a

product in order to fraudulently receive a refund. Such scamming is anticipated by users (Van Hout and Bingham (2014)) and a motivation for many sites requiring users to leave feedback or instituting mediation systems. As with hacks, scams are identified through user claims.

An additional adverse event is the denial-of-service (DoS) attack. In this attack, attackers disrupt a service to prevent legitimate users from accessing it. For example, attackers might overwhelm a website with information requests so that legitimate visitors are unable to load it. DoS attacks occur frequently on the ecosystem however they are difficult to study because many DNMs provided such irregular service that it is often difficult for observers to determine if a DNM is offline due to a bug or because of a deliberate DoS attack.

3.2.1 Impact on Ecosystem

Each of these events has the ability to affect the resilience of the ecosystem even though it does not target the ecosystem as a whole. This is because they hinder the ability of the ecosystem to perform its primary function of facilitating trade. The ways in which the ecosystem is impacted are as follows:

- **DNM Shut Down**: Whether one DNM is closed or many, removing DNMs reduces the immediate overall size of the ecosystem and can do so dramatically if many accounts are lost. In some instanses, the removal of one or more DNM has led to an increase in population size (for example when *Silk Road* was closed the remaining markets saw a large increase in users Christin (2013). This is in part due to the displacement of existing users and due to new ones joining. The mass displacement of users, as occurred after the closure of the DNM *Alphabay* (Greenberg (2017)), can slow down the service provided by other DNMs that are unable to cope with sudden influxes of users. Finally, the removal of DNMs may reduce the availability of certain products and, therefore, the overall service provided.
- **User Warning**: law enforcement targeting users with warnings may deter users from participating in the ecosystem thus reducing the overall population.
- **Arrest**: Arrests of vendors reduce the amount of trade on the ecosystem, especially if they have a large number of customers.
- **Parcel Seizure**: When parcels are seized during transit this prevents the ecosystem from fulfilling its main objective - to facilitate more convenient trade. This is because, even if vendors reimburse buyers or reship the product, the overall costs and the amount of time spent on purchases increases. Further, to avoid the seizure

TABLE 3.1: Specific Impacts on Ecosystem Attributed to Each Possible Event

Event	Impact on Ecosystem
DNM Shut Down	Reduce Trade Reduce Population Reduce Cashflow Reduce Convenience
User Warning	Reduce Population
Arrest	Reduce Population
Parcel Seizure	Reduce Convenience
Closure	Reduce Trade Reduce Convenience
Exit Scam	Reduce Trade Reduce Population Reduce Cashflow Reduce Convenience
Hack	Reduce Cashflow
Scam	Reduce Convenience

of their parcels, vendors must invest in more stealth strategies which also reduces the convenience of the service.

- **Closure**: When DNMs are closed the ecosystem becomes smaller and trade is reduced. As with a shut down, this can also reduce service.
- **Exit Scam**: Exit scams also reduce the size of the ecosystem but, after these incidences, less coin is likely to be reinvested into the system as it is stolen in the scam. In addition the population of the ecosystem and overall trade and service is reduced.
- **Hack**: Hacks can cause DNMs to close temporarily (DeepDotWeb (2014d)) reducing the available service to users. Additionally, if coin is stolen then these events also reduce the amount of cashflow across the ecosystem.
- **Scam**: When scams occur, as with seized packages, users are not receiving a more convenient and safer form of trade therefore the presence of scams on the ecosystem affects its resilience.

A summary of the possible impacts of each event is given in Table.3.1. This gives 4 impacts on the ecosystem: a reduction in trade (amount and value of sales taking place), population (number of users engaged in the ecosystem), cashflow (value of coin available for spending) and convenience (availability, diversity and price of products and services) which may combine to diminish the capacity of the ecosystem.

It is likely that the impacts are not independent and have the ability to influence each other. For instance, a reduction in the population of the ecosystem will also reduce the overall cashflow as there are fewer buyers. This, in turn, would reduce the quantity of

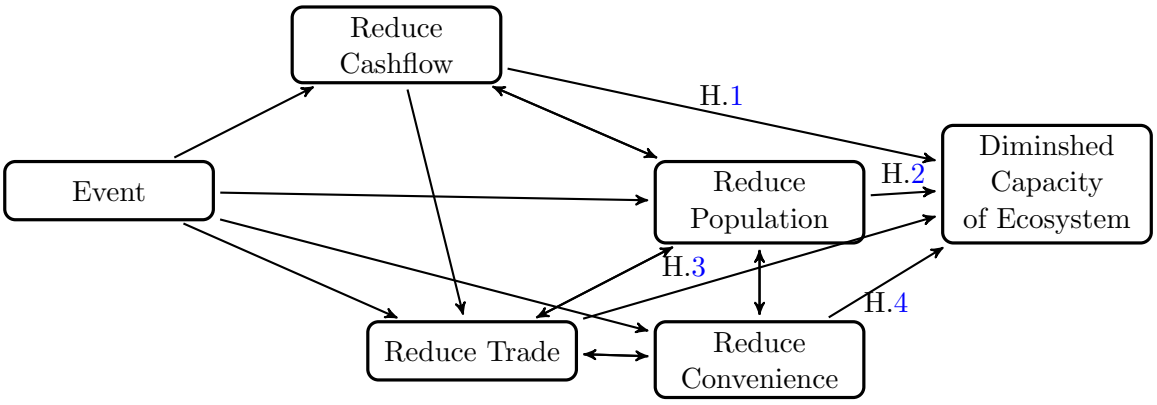

FIGURE 3.1: **Hypothesis Model for the Impacts of Events on Ecosystem**

trade. The way that the different impacts interact to affect the capacity of the ecosystem is presented in Figure 3.1 as a model for how events impact the ecosystem.

To understand if this model is accurate, the following hypotheses will be tested.

3.2.1.1 Ecosystem Hypotheses

If a reduction in the amount of coin available for spending on the ecosystem diminishes the capacity of the ecosystem then events that remove more coin should have a bigger impact.

Hypothesis 1. As Operation Onymous removed $1 million and the Evolution Exit Scam removed between $12 million and $34 million, the Evolution Exit Scam had a bigger impact on the ecosystem than Operation Onymous.

If a reduction in the population diminishes the capacity of the ecosystem then events that reduce the population should have a bigger impact than a comparable event that does not reduce the population.

Hypothesis 2. As arrests reduce the population but package seizures do not, heightened periods of arrests should diminish the capacity of the ecosystem more than heightened periods of parcel seizures.

Some law enforcement interventions have had a direct impact on the trade occurring on the ecosystem by removing the DNMs that facilitate it. These have also removed substantial amounts of coin from the system. Others, such as Operation Hyperion, have not had such impacts and, if the reduction of these trade related elements of the ecosystem affect its resilience, then we would expect less activity after events such as Operation Onymous and Operation Bayonet than Operation Hyperion.

Hypothesis 3. Operation Onymous and Operation Bayonet had a greater impact on the ecosystem than Operation Hyperion.

If the convenience of the ecosystem is a determinant for its capacity then we would expect to see a lesser capacity during periods when convenience is low, e.g. when there are less products available, more scams reported, and/or higher prices.

Hypothesis 4. The capacity of the ecosystem is perceived to be diminished when the convenience provided is also perceived to be low.

Figure 3.1 argues that a reduction in the population both affects and is affected by the other variables. Whether or not, and how, this might be the case is discussed in the next section which explores the impact of events on individual users.

3.2.2 Impact on Users

The ways in which each event could impact on individual users are:

- **DNM Shut Down**: When law enforcement shuts down a DNM the immediate impact is the loss of vendor and buyer accounts which can no longer be accessed. Not only are users denied the ability to access their accounts and make future trades or attend to existing ones, they also lose any of their coin kept in escrow on the DNM. They will also lose the reputation they have accumulated for that account if they are unable to transfer it to a new account on a different DNM.
  In addition, users are now aware that law enforcement has access to any personally identifiable information that they might have shared privately on the messaging services of the DNM. This may cause them to want to stop trading altogether in fear of providing additional evidence for any law enforcement investigation opened on them. Indeed, when interviewed, many sellers on Silk Road were aware of the amount of information stored on the site's servers that could be used to deanonymise them (Van Hout and Bingham (2014)). A residual impact of law enforcement interventions is that users not active on the DNM shut down may be deterred from continuing to trade because they are afraid of future law enforcement interventions. Similarly, users may become wary of the instability of the overall ecosystem and the fact that DNMs can disappear without prior warning, increasing the risk of trade.
  Another possible impact is the potential loss of trading partners, this could be because a user's trading partners leave the DNM or because they are unable to reconnect with their partners on another site. In either instance, this may still hinder a user's ability to trade on the ecosystem.
- **User Warning**: When users are approached by law enforcement and warned they might be arrested, they may become more worried about future law enforcement interventions and being arrested.

- **Arrest**: When a user is arrested any assets related to their illegal activities can be seized and their account shut down. This clearly will restrict on that particular user's ability to continue trading but, in addition, other users in the ecosystem may also be affected. Anyone trading with the user arrested will lose a trading partner and may lose payment or product that has yet to be dispatched. Publication of the arrest may also increase fear of law enforcement.
- **Parcel Seizure**: If a parcel is seized before it reaches its destination it could result in the buyer having an unsuccessful trade because they are unable to get a (full) refund and/or the vendor receiving bad feedback because of the failed transaction. Alternatively, the vendor may choose to resend the product or refund the buyer in which case they have reduced their profits for that purchase.
  Parcel seizures can also heighten user awareness of law enforcement, particularly for users who are fearful of being under investigation or for users who received a "Love Letter" which explicitly notifies them of the parcel seizure. This may deter them from continuing to trade causing others to lose their trading partners.
- **Closure**: When a DNM is closed, users have the opportunity to make plans with their trading partners, remove any coin they have in escrow and move to another DNM. However, they may still lose trading partners if their customers cannot follow them to another site/vendors do not create accounts on a different site and the account that they created and built a reputation upon is lost. Further, if this event were to happen with frequency it may cause users to fear that the ecosystem as a whole is unreliable or unstable.
- **Exit Scam**: As with the DNM Shut Down, an exit scam would cause users to lose their accounts and, therefore, any coin held on the site. They may also lose their trading partners if customers/vendors quit the ecosystem or are unable to return after an exit scam. However, an exit scam would not necessarily cause users to be fearful of deanonymisation as the servers containing their information is not taken by law enforcement. Instead, they may become wary of returning to the ecosystem if they are concerned that more exit scams will take place.
- **Hack**: When a DNM is hacked some users may have their coin stolen. Some DNMs, when this has occurred, have tried to pay back the amount each user lost (DeepDotWeb (2014d)), however this is not the case for all. In addition to stealing coin, hackers can steal information from servers and this may make users afraid of law enforcement investigations.
- **Scam**: If a vendor deliberately does not send a product, or sends a product of lower quality than expected, or a buyer refuses to pay for a product that has arrived as expected, then the trading partner will lose profit on this transaction.

TABLE 3.2: Specific Impacts on Users Attributed to Each Possible Event

Event	Impact on Affected Users	Impact on Unaffected Users
DNM Shut Down	Lose Account Lose Coin Lose Trading Partners Fear of Law Enforcement Fear of Instability	Lose Trading Partners Fear of Law Enforcement Fear of Instability
User Warning	Fear of Law Enforcement	Fear of Law Enforcement
Arrest	Lose Coin Lose Product Fear of Law Enforcement	Lose Trading Partners Fear of Law Enforcement
Parcel Seizure	Lose Coin Lose Product Lose Reputation Fear of Law Enforcement	Fear of Law Enforcement
Closure	Lose Account Lose Trading Partner Fear of Instability	Fear of Instability
Exit Scam	Lose Account Lose Coin Lose Trading Partners Fear of Instability	Fear of Instability
Hack	Lose Coin Fear of Law Enforcement Fear of Instability	Fear of Law Enforcement Fear of Instability
Scam	Lose Coin Lose Product	

The proposed impacts of each event are summarised in Table. 3.2. The impacts are considered in the context of users directly affected by an event and those who are indirectly affected, e.g. are present on the ecosystem but do not have an account on the DNM shut or are themselves arrested, etc. These impacts are the risk factors that may lead to a user becoming maladjusted.

From this categorisation, it can be seen that affected users can be impacted by an event because they lose their account, their coin, their reputation, their product and/or their trading partners or because they feel fear of law enforcement and/or instability and unaffected users can be impacted because they can lose their trading partners and feel fear of law enforcement and/or instability.

The way that these impacts interact is described in Figure 3.2. And, the following hypotheses are given to evaluate if the events impact the population in the ways described.

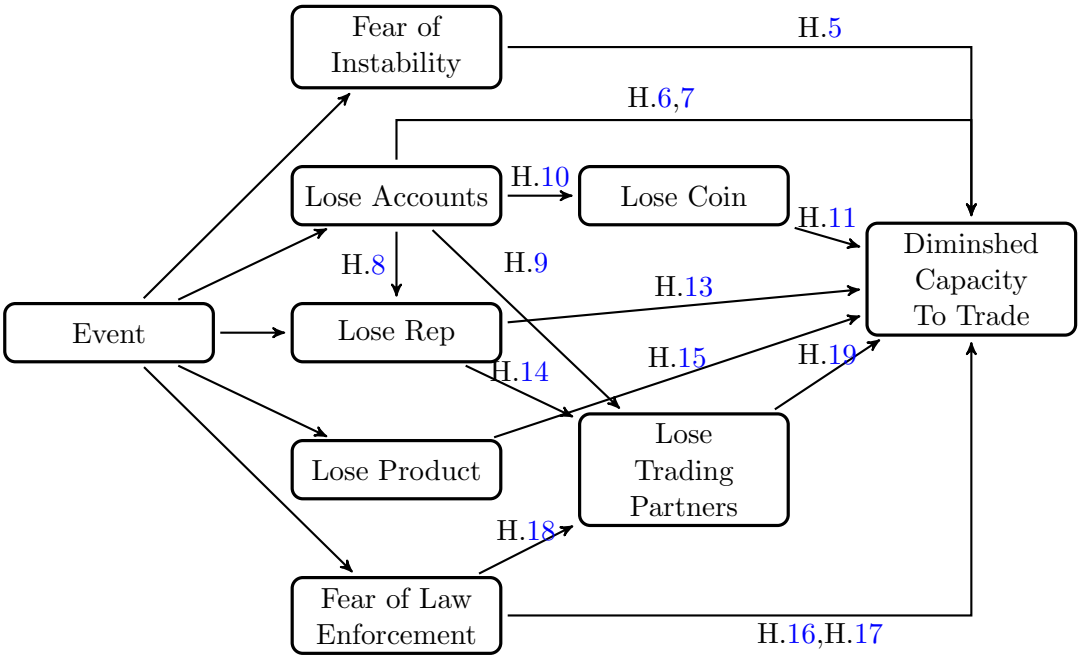

FIGURE 3.2: **Model for the Impacts of Events on Users**

3.2.2.1 User Hypotheses

In addition to a service being present, users must feel that it will remain present and be reliable in order to trade on it and, therefore risk coin or information that they have stored in the system. Users who have seen multiple DNMs close may be more aware of the instability of the ecosystem and so more likely to leave it. In particular in Operation Bayonet, the DNM that many users moved to (*Hansa*) after the closure of *Alphabay* was also closed. This event therefore caused more instability than other DNM closures and, as such, it should be expected that this event had a bigger impact on its affected population.

Hypothesis 5. The more DNMs that a user has witnessed close, the more likely they are to leave the ecosystem after an event.

Figure 3.2 argues that losing an account diminishes the capacity of users to trade. If this is the case then it is expected that users who lose accounts are more likely to stop trading. Not only this, but users who have less accounts would be expected to be more greatly affected and therefore more likely to stop trading.

Hypothesis 6. If losing an account is a risk factor which makes users less resilient then users who lose accounts are more likely to stop trading on the ecosystem.

Hypothesis 7. If the impact of losing an account more greatly affects users with fewer accounts then the more accounts a user owns, the more likely they are to continue trading after losing an account.

Losing an account is not only predicted to directly diminish a user's capacity to trade but also to indirectly do so by reducing a user's reputation, the amount of coin they have and their number of trading partners.

Hypothesis 8. If losing an account has an impact on user reputation then users will be unable to retain their reputation after an account they are active on closes.

Hypothesis 9. If losing an account has an impact on maintaining relationships then users will be unable to find their trading partners after their account is closed.

Hypothesis 10. If losing an account has an impact on the amount of money users have then users will record financial losses after events that close their accounts.

If losing coin stored on a DNM has an impact on the resilience of users, for example by preventing them from being able to buy products or stock to sell or by reducing the profitability of their participation, then events that result in a loss of coin will have a bigger impact than similar events that do not. This would be the case when comparing exit scams or closures by law enforcement to amicable closures because, in the former instances, users do not get enough warning to empty their accounts and are more likely to lose money. Further, when events affect some users by taking their coin but not other users, then those users who lose their coin should more likely to leave the ecosystem.

Hypothesis 11. When a DNM is closed by law enforcement or in an exit scam, more users leave the population than after a DNM is closed amicably by its owners.

In order to understand whether it is the loss of an account, or the loss of coin held on an account, or both which most diminishes a user's capacity to trade, an event in which users lost their accounts and coin (the Evolution Exit Scam) was compared to an event in which users just lost money (the Silk Road 2.0 hack). If losing an account, rather than the coin held on the account, diminishes a user's capacity to trade then it should be expected that proportionately more users stop trading after the Evolution Exit Scam.

Hypothesis 12. When an event results in the closure of accounts and the loss of coin more users leave the population than after an event in which users just lose coin.

Another attribute that users hold which may affect their capacity to trade after an event is their reputation. Figure 3.2 predicts that losing reputation diminishes a user's capacity to trade, if this is true then users with greater reputation should be more capable of continuing to trade after an event in which reputation is lost.

Hypothesis 13. If losing reputation is a risk factor which makes users less resilient then users with lower reputations are more likely to stop trading on after events.

It has already been shown that vendors reduce the prices of their products when their reputation falls (Hardy and Norgaard (2016)). It is therefore presumed that a loss in reputation results in a loss in trading partners.

Hypothesis 14. After an event, vendors who lose reputation lose trading partners.

Some events may result in a buyer losing a product they had purchased or a vendor losing a product they intended to sell.

Hypothesis 15. If losing a product because of an event impacts a buyer's ability to continue trading then buyers who demonstrate they have lost a product will be more likely to decide to stop trading.

Exit scams and the closure of DNMs by law enforcement have similar impacts on users: they both result in the loss of accounts, trading partners, and coin held in escrow. However, when a DNM is shut down by law enforcement affected users may additionally be worried about law enforcement using the information they stole to arrest them. As a result, it is expected that a greater proportion of users leave the ecosystem after a law enforcement operation like Operation Onymous than after an exit scam of similar scale, such as that of the DNM Evolution.

Hypothesis 16. A greater proportion of users left the ecosystem after Operation Onymous than after the Evolution Exit Scam.

If the fear of being arrested is sufficient to deter users from participating in the ecosystem at all then events, such as Operation Hyperion, that do not incur any of the other impacts brought about by different events should lead to a reduction in the population of the ecosystem.

Hypothesis 17. Operation Hyperion led to a reduction in the size of the ecosystem population.

This fear of law enforcement may be manifested in specific users, as opposed to the ecosystem as a whole. In these instances, rather than leaving the ecosystem, users may instead stop trading with particular actors.

Hypothesis 18. If an increased fear of law enforcement has the capacity to cause users to lose their trading partners then users demonstrate they have stopped trading with other users for this reason.

Finally, users establish trading relationships with each other over time. If these relationships are valuable and user specific, i.e. cannot be easily transferred to another,

comparable user, then losing a trading partner would diminish a user's capacity to trade. If this is the case we would expect users who lose trading partners to be less likely to continue trading.

Hypothesis 19. If losing a trading partner is a risk factor which makes users less resilient then users who lose trading partners are more likely to stop trading on the ecosystem.

Chapter 4

Data

This section discusses three datasets which have been used to answer the previously outlined research question and hypotheses.

4.1 Gwern Branwen Dataset

The independent researcher Gwern Branwern has collated a large, public dataset containing downloads from several DNMs (Branwen et al. (2015))[1]. Between July 2013 and July 2015 87 DNMs and over 37 related forums were crawled and mirrored (downloaded). Of these, 78 DNMs were drug related or focused:

1776, Abraxas, Agape, Agora, Alpaca, Alphabay, Amazondark, Anarchia, Andromeda, Area 51, Atlantis, Black Goblin, Black Market Reloaded, Black Services Market, Bloomsfied, Blue Sky, Breaking Bad, Cannabis Road, Cannabis Road 2, Cannabis Road 3, Cantina, Cloud 9, Cryptomarket, Dark Net Heroes, Darkbay, Darklist, Deep Bay, Deepzon, Diabolus, Dogeroad, Dreammarket, Drugslist, Eastindia Company, Evolution, Freebay, Freedom Market, Freemarket, Grey Road, Havana, Haven, Horizon, Hydra, Ironclad, Kiss, Middlearth, Mr Nice Guy 2, Nucleus, Onionshop, Outlaw Market, Oxygen, Panacea, Pandora, Pigeon Market, Pirate Market, Poseidon, Sheep, Silk Road, Silk Road 2.0, Silk Road Reloaded, Silk Street, Simply Bear, the Black Box, the Majestic Garden, the Marketplace, the Real Deal, Tochka, Tom, Topix 2, Tor Escrow, Tor Market, Torbay, Torbazaar, Tortuga 2, Underground Market, Utopia, Vault 43, White Rabbit, Zanzibar Spice

[1] Other, similar datasets (containing scrapes from the public facing pages of DNMs) have been used by other researchers. Some of these are more complete than the dataset used and were available for use by the author. However, this dataset was used primarily because of its availability and perceived usefulness at the beginning of the project. It was then not feasible, due to the time constraints of the funding for this PhD, to incorporate other datasets when they became available.

This data has been used to conduct several dozen studies and research projects (Branwen et al. (2015)). The research using this dataset has built conclusions on the nature of forum discussions (Munksgaard and Demant (2016)), the impact of law enforcement interventions (Décary-Hétu and Giommoni (2016)), and even the validity of other research (Munksgaard et al. (2016)). However, this data has not been explicitly validated and it has even been claimed that it is impossible to determine how incomplete the data is as a representation of the ecosystem during the measurement period (Décary-Hétu and Giommoni (2016)). The lack of measurements on the data means that results drawn from it are potentially not as reliable as other, peer-reviewed datasets (Dolliver (2015b)). As such, a number of different metrics were employed to assess the quality of data collected on each DNM and data that was determined to be of too low quality was discarded. Section. 4.1.1 describes how information was first extracted from the scrapes and Section. 4.1.2 describes how it was validated.

4.1.1 Information Extraction

Each mirror was formatted in html and parsed using BeautifulSoup and xTree for Python 2.7. Information was extracted from the profile and listing pages on the site. On the few occasions where profile or listing pages were not available (*Agape*, *Alphabay*, *Black Services Market*, *Breaking Bad*, *Dream Market*, *Onion Shop*, *Tochka*, *Torbay*, *Torbazaar*, *Tortuga2*, *White Rabbit*), as much information as possible was collected elsewhere, such as from the index pages or review pages. For each vendor, at least the name and date of the scrape was extracted as well as some combination of the following:

- profile description,
- PGP key,
- reputation value,
- registration date,
- the date the vendor was last active,
- reviews,
- product listings,
- sales statistics,
- any other sites on which they claim to be active.

For each product listing, at least the title of the product and the date of the scrape was extracted, as well as any of the following information if it was available:

- the product description,
- price,
- vendor,

- vendor and product reputation values,
- product feedback,
- sales statistics,
- where the product is shipped from and to,
- the shipping costs,
- the product category,
- the quantity of stock still available,
- the maximum order size.

447,084 individual profile pages and 6,605,053 listing pages were parsed, with an average of 6,985 and 103,203 per marketplace, respectively. This totalled 12.5 GB of data. Some of the scrapes contained pages that were empty, displayed a log in screen only, or an error code. 288,495 pages contained no information across all the DNMs, with an average of 4,507 per DNM, in total this was 82.7MB of data. Refer to Table. 4.1 for more specificity.

4.1.2 Data Validation

To begin with 14 DNMs were discarded because the data collected did not give an accurate picture of each DNM, this was either because too little information was contained in the scrape or because there was an issue with data collection. Where there was only one scrape for a DNM, it was discarded as it has been shown that more scrapes lead to greater coverage (Soska and Christin (2015)). These DNMs were:

- **Agape** the scrape did not contain any vendor information.
- **Atlantis** there was only one scrape collected on a date after the closure of the DNM.
- **Black Goblin** one scrape was collected and, at this point, there were no listings advertised on the DNM.
- **Black Market Reloaded** the scrape captured only a small section of the DNM at the end of its lifetime - there were only two scrapes and they contained index pages but no profile pages or listing pages.
- **Cannabis Road** the scrape captured no information - there were no profile or listing pages and the index page was logged out.
- **Deep Bay** the web pages did not display relevant information and, instead, redirected the visitor to another site.
- **Freedom Market** the scrape only collected one user profile - the one created in order to collect information, and no listing pages.

TABLE 4.1: Sizes of Datasets

Marketplace	Number of Vendor Profile Pages	Number of Listing Pages	Number of Empty Pages	Date of First Scrape	Date of Last Scrape	Number of Scrapes
1776	572	1929	0	08/05/2014	20/09/2014	26
Abraxas	17795	315506	0	16/12/2014	05/07/2015	95
Agora	117191	2386300	100209	01/01/2014	07/07/2015	204
Alpaca	3390	5282	0	24/04/2014	07/11/2014	34
Alphabay	7437	34413	7690	22/12/2014	04/07/2015	72
Amazondark	268	1209	0	11/06/2015	04/07/2015	10
Anarchia	770	1880	0	07/05/2015	05/07/2015	19
Andromeda	2054	9061	399	12/04/2014	17/11/2014	37
Area 51	3372	15345	0	22/06/2014	23/01/2015	60
Black Services Market	310	942	78	26/12/2013	14/02/2014	7
Bloomsfied	56	73	0	19/02/2015	29/03/2015	20
Blue Sky	188	1470	0	06/01/2014	28/09/2014	10
Breaking Bad	1	7	0	01/02/2014	02/02/2014	2
Cannabis Road 2	1101	6652	0	03/04/2014	05/08/2014	22
Cannabis Road 3	914	522	148	05/10/2014	31/10/2014	8
Cantina	27	38	0	20/01/2014	07/02/2014	9
Cloud 9	16418	92408	58	11/02/2014	01/11/2014	37
Cryptomarket	13361	17341	1	19/02/2015	06/07/2015	38
Dark Net Heroes	272	1237	1	30/05/2015	04/07/2015	12
Darkbay	1188	1908	85	30/01/2014	13/05/2014	13
Darklist	506	8	0	13/02/2014	04/04/2014	4
Deepzon	49	262	0	22/05/2014	05/07/2014	7
Diabolus	11303	51041	1	17/10/2014	05/07/2015	112
Dogeroad	117	408	0	20/01/2014	28/02/2014	7
Dreammarket	6257	73621	3	09/01/2014	05/07/2014	151
Drugslist	299	2204	301	09/01/2014	12/05/2014	16
Eastindia Company	1728	8911	1	28/04/2015	05/07/2015	23
Evolution	19595	502545	26188	21/01/2014	17/03/2015	115
Freebay	975	2567	0	01/01/2014	21/02/2014	7
Freemarket	104	4729	1	15/01/2015	26/02/2015	40
Greyroad	11	33	0	15/01/2014	25/01/2014	2
Haven	432	3321	0	10/05/2015	05/06/2015	8
Horizon	40	176	0	28/06/2015	04/07/2015	4
Hydra	4276	48465	1	03/04/2014	27/10/2014	36
Ironclad	562	3605	1	17/03/2015	24/03/2015	6
Kiss	582	3604	0	19/02/2015	15/05/2015	37
Middlearth	7000	34257	355	23/06/2014	05/07/2015	119
Nucleus	22254	47854	7	24/10/2014	07/07/2015	117
Onionshop	258	320	0	21/05/2014	27/10/2014	25
Outlaw Market	84	9217	5	09/01/2014	05/07/2015	114
Oxygen	2504	50680	144	26/04/2015	05/07/2015	22
Panacea	558	6724	0	28/10/2014	12/02/2015	54
Pandora	33397	629801	42176	25/12/2013	05/11/2014	105
Pigeon Market	15	100	0	21/04/2014	24/04/2014	2
Pirate Market	3946	11387	10814	26/12/2013	21/09/2014	60
Poseidon	136	3688	0	05/06/2015	04/07/2015	11
Sheep	1993	6687	24	17/11/2013	03/12/2013	1
Silk Road 2	119163	2156343	84276	20/12/2013	06/11/2014	57
Silk Road Reloaded	301	3543	517	18/01/2015	05/07/2015	56
Silk Street	65	179	0	24/04/2014	27/07/2014	13
the Majestic Garden	224	524	0	24/04/2014	15/09/2014	15
the Marketplace	12379	17793	231	03/01/2014	09/11/2014	61
the Real Deal	839	8334	3100	16/04/2015	05/07/2015	30
Tochka	760	495	0	05/02/2015	04/07/2015	45
Tom	889	6256	11389	05/05/2014	17/12/2014	42
Topix 2	4197	985	1	28/09/2014	05/11/2014	33
Tor Escrow	284	1408	1	15/01/2014	20/04/2014	10
Tor Market	863	5110	0	06/12/2013	14/12/2013	3
Torbay	198	87	0	01/01/2014	12/04/2014	10
Torbazaar	582	1901	1	02/02/2014	06/11/2014	40
Tortuga2	32	268	0	23/04/2014	09/06/2014	8
Underground Market	94	759	1	24/04/2014	27/08/2014	18
Utopia	3	3	40	11/02/2014	23/02/2014	2
White Rabbit	545	1327	247	15/12/2013	14/04/2014	12

- **Havana** the web pages did not display relevant information and, instead, displayed a wait screen image.
- **Mr Nice Guy 2** the files were corrupted and so no information could be extracted.
- **Sheep** only one scrape was taken.
- **Silk Road 1** none of the scrapes of Silk Road 1 included in the dataset were collected by the same researcher.
- **Simply Bear** only 2 vendor profiles were collected in the scrape, 1 of which was an administrator, and no listing pages were collected.
- **The Black Box** this DNM facilitates encrypted, private sales between vendors and buyers, there are no profile or listing pages in the scrape.

- **Utopia** only one scrape was taken.
- **Vault 43** there was no vendor or listing information captured in the scrape.
- **Zanzibar Spice** the only vendors in the scrape were admin or test vendors, and the only listings were test listings.

For the remaining DNMs, their scrapes could be incomplete for three reasons: firstly because each individual scrape was not able to collect all of the available information on the DNM, this could be because the DNM went offline during the scrape (Van Buskirk et al. (2015)) or there was an issue with the Tor network (Branwen et al. (2015)); the scraping tool was blocked by a CAPTCHA or other technology; the scraping tool was unable to access all of the pages on the DNM, for instance because the hyper-links were not easily accessible (Munksgaard et al. (2016)); or, simply because the DNM was too large to scrape in the time period given (Christin (2013)). Secondly because the scrapes were collected at too large time intervals and so information which was only available temporarily is lost. And, thirdly, because not all of the DNMs active during the measurement period were scraped. The researcher who collected the scrapes has explicitly stated that each of the scrapes is an underestimate of the information that was available (Branwen et al. (2015)), but it is still not understood to what extent the data recorded is an underestimate and whether or not it can still be used to make meaningful conclusions about the nature of DNMs.

It was first determined which DNMs could not be used because the scrapes collected on them did not seem to give an accurate image of the DNM. As opposed to looking for the most complete scrapes only, the data for each DNM was evaluated to see if it represented the change in size of the DNM over time even if only a fraction of the DNM was captured in each scrape. This is because, if the dataset consistently underestimates the amount of data held on a site but does so proportionately as the site changes in size then, the dataset may still be used to understand the growth of the ecosystem over time. However, if the pattern of growth recorded by the dataset diverges from reality then sudden changes in the observed population caused by errors in data collection may be misattributed to events leading to erroneous conclusions.

To estimate the expected growth of each DNM and the amount of data that was not collected, a number of different metrics were identified. These metrics provided comparative estimates on the number of vendors and listings collected as they are key variables for the analysis in the study. Appendix A describes in greater detail the exact amount of information taken from each DNM for the purposes of validation. These figures were also found through a second dataset collected from the search engine Grams.

The first metric used the advertised number of vendors and the advertised number of products. These were found on 3 and 42 DNMs respectively. There is no way of

TABLE 4.2: Kolmogorov-Smirnov Test Results on the CDF of the Advertised and Recorded Numbers of Listings

DNM	Test Statisitc	p-Value
1776	0.120	0.990
Abraxas	0.065	0.988
Agora	0.078	0.584
Alpaca	0.161	0.778
Alphabay	0.167	0.246
Amazondark	0.100	1.00
Anarchia	0.053	1.00
Andromeda	0.162	0.676
Area51	0.050	1.00
Black Services Market	0.857	0.004
Bloomsfield	0.053	1.00
Blue Sky	0.800	0.001
Cannabis Road 2	0.091	1.00
Cannabis Road 3	0.5	0.188
Cantina	0.333	0.810
Cloud 9	0.100	0.997
Cryptomarket	0.438	0.003
Darkbay	0.167	0.991
Diabolus	0.091	1.00
Dogeroad	0.143	1.00
Dreammarket	0.027	1.00
Drugslist	0.214	0.862
Eastindia Company	0.045	1.00
Evolution	0.0818	0.840
Freebay	0.143	1.00
Haven	0.125	1.00
Horizon	0.0	1.0
Ironclad	0.167	1.00
Kiss	0.100	0.997
Middleearth	0.280	0.000142
Nucleus	0.261	0.000603
Panacea	0.0556	1.00
Pirate Market	0.220	0.0982
Poseidon	0.0909	1.00
Silk Road 2.0	0.137	0.702
Silk Street	0.0769	1.00
the Real Deal	0.105	1.00
Tochka	0.333	0.0277
Topix 2	0.0303	1.00
Tor Escrow	0.667	0.320
Tor Market	0.5	0.844
Torbay	0.100	1.00
Underground Market	0.0625	1.00

verifying if these figures are accurate for each DNM, and it has been observed that the advertised number can overestimate the actual number of available products (Dolliver (2015b)). It is known, for example, that some marketplaces would report the number of listings that had ever been available, rather than just the number that was available at any one time. Even if they are not perfectly accurate, however, the motivation to overestimate sales numbers or community sizes would likely remain consistent over time and therefore reflect periods of the site's growth. Additionally, where the advertised number of products was calculated by summing the advertised number of products for each category, the value used for comparison may be too large if products could be listed in more than one category and so are counted twice. This is another reason why the figures advertised are unreliable for quantifying the amount of data missing.

These fluctuations in size were compared using the Cumulative Distribution Function (CDF) of both the advertised and observed listings. The number of listings were added cumulatively to more clearly see the rate of growth. For each DNM, the CDFs of the number of listings collected over time was compared to the number of listings advertised over time using the Kolmogorov–Smirnov test and the results are presented in Table. 4.2.

Where the p–value was less than the significance level of 5%, i.e. $p < 0.95$ for a DNM, it was determined that the scrapes could not be used to measure the activity on the DNM. The CDFs for these DNMs were inspected to see if there were subsections of the scrape which could be used. The test failed for *Agora, Alpaca, Alphabay, Andromeda, Black Services Market, Blue Sky, Cannabis Road 3, Cantina, Cryptomarket, Drugslist, Evolution, Hydra, Middle Earth, Nucleus, Pirate Market, Silk Road 2.0, the Market Place, Tochka, Tor Escrow,* and *Tor Market* from the dataset. The CDF plots for each DNM can be seen in Appendix G.

The 3 DNMs that displayed the number of vendors on their site were *Cloud 9*, *Outlaw Market*, and *Panacea*. Both *Cloud 9* and *Panacea* listed the number of vendors that were shipping products from different destinations and, for each DNM, the advertised number of vendors was taken to be the sum total of these values. Outlaw Market gave users an opportunity to filter listings by vendor and so displayed a list of the vendors on the DNM, the advertised number of vendors was taken to be the length of this list.

As with the comparison of listings, the CDF of the observed vendor population and the recorded vendor population was compared. Some researchers have tried to account for the incompleteness of scrapes by assuming vendors to be present from the first time they appear in a scrape until the last time, even if their profile is not collected in scrapes taken between these two dates. To see if this smoothing mechanism increases the accuracy, the CDF of the observed vendor population measured in this manner was also compared to the advertised population CDF. The Kolmorgov–Smirnov test confirmed the similarity of the CDFs for both *Cloud 9* and *Pancea* with test statistics and p-values of 0.0741, 0.998 and 0.0667, 1.00 respectively. The CDF plots are displayed in Figure 4.1. For *Cloud 9*, the similarity between the CDF of the advertised vendor population and smoothed, observed population was slightly weaker than with the observed population as the p-value of the KS test in this incidence was 0.997. For *Panacea*, the KS test produced the exact same result.

In contrast, the test failed for *Outlaw Market*. This was also determined by the KS test which was 0.167 with p-value 0.945.

The second metric utilised a second dataset collected from the Grams site, a Dark Net search engine which allowed buyers to find vendors and products across different DNMs. This dataset contained 419 scrapes across 418 dates from 9 June 2014 until 17 April 2016. It contained information on 18 relevant DNMs in the form of '.csv' files which recorded listings along with their vendor, price, description, the time they were added, and the location they ship from. It was also collected by independent researcher Gwern Branwen.

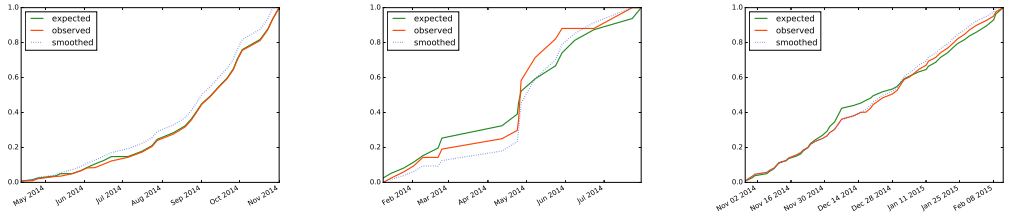

(a) **CDF plot of Cloud 9** test statistic 0.0667, p-value 1.00

(b) **CDF plot of Outlaw Market** test statistic 0.167, p-value 0.945

(c) **CDF plot of Panacea** test statistic 0.0741, p-value 0.998

FIGURE 4.1: **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Vendors as well as the CDF Plot of the Smoothed Vendor Population.

TABLE 4.3: Kolmogorov-Smirnov Test Results on the CDF of the Number of Vendors Found in the Grams and Branwen Datasets

DNM	Observed		Smoothed		Number of Scrapes
	Test Statisitc	p-Value	Test Statisitc	p-Value	
1776	0.00	1.00	1.00	0.289	1
Abraxas	0.0476	1.00	0.0952	1.00	21
Agora	0.0286	1.00	0.0286	1.00	105
Alpaca	0.167	1.00	0.167	1.00	6
Alphabay	0.174	0.842	0.217	0.593	23
Andromeda	0.231	0.828	0.385	0.226	23
Cloud 9	0.0769	1.00	0.154	0.995	13
Dreammarket	–	–	–	–	0
Evolution	–	–	–	–	0
Haven	0.500	0.844	0.500	0.844	2
Nucleus	0.125	0.893	0.075	1.00	40
Outlaw Market	–	–	–	–	0
Oxygen	0.143	1.00	0.143	1.00	7
Pandora	0.100	1.00	0.100	1.00	10
Pirate Market	0.143	0.973	0.0952	1.00	21
Silk Road 2.0	0.0870	1.00	0.0870	1.00	23
the Real Deal	–	–	–	–	0
Tom	0.143	1.00	0.143	1.00	7

The DNMs found in both datasets were *1776, Abraxas, Agora, Alpaca, Alphabay, Andromeda, Cloud 9, Dreammarket, Evolution, Haven, Nucleus, Outlaw Market, Oxygen, Pandora, Pirate Market, Silk Road 2.0, the Real Deal, Tom.*

The vendor population captured in the Branwen scrape was compared to the vendor population collected from Grams. To decide which DNMs to discard as a result of this method, the CDF of the two vendor populations were compared using the Kolmogorov-Smirnov test.

This analysis was repeated with a modified observed vendor population created by identifying the full lifetime of each vendor present on a DNM in the Branwen dataset and then assuming they are present for the totality of this lifetime, even if their profile does not appear in each scrape (Soska and Christin (2015)). The analysis was repeated to determine whether or not this technique can be employed to improve completeness of a dataset. The results of the Kolmogorov–Smirnov Test are presented in Table. 4.3.

The KS test failed for *Alphabay*, *Andromeda*, *Haven*, and *Nucleus* when comparing the observed vendor populations. However, for *Nucleus*, the KS test passed when comparing the population observed on Grams to the smoothed population. In this instance, measuring the vendor population by counting vendors as present from their first appearance until their last may provide a more complete picture of the DNM. However, the KS-test produced a worse score on this comparison for *Alphabay*, *Andromeda*, and *Cloud 9*. This may be because the Grams dataset wasn't smoothed (whether or not it was is unknown) and so smoothing the data made it less similar to the trend found in Grams.

Both this test on the vendor population and the previous test on listing information failed for *Alphabay*, and *Andromeda*. As such, they are removed from the dataset. For *Alpaca*, *Agora*, *Nucleus*, *Pirate Market*, and *Silk Road 2.0*, the test passed for the vendor population but failed for listings, therefore these DNMs are kept in the dataset when analysing the vendor population but not when producing results on listings. This test failed for *1776* and *Haven* but passed on the comparison of their listings. These DNMs were kept in the dataset because there was such a small intersection between the Branwen and Grams datasets in their case. This means that the second test analysed substantially less of the dataset and provided a less trustworthy conclusion.

For the DNM Evolution, the CDF clearly shows that, from September 2014, the increase in observed and expected size becomes much more similar. Indeed, the KS-test on data after this date only produces a test statistic of 0.0556 with p-value 1.00. As such, this subsection of Evolution's scrapes is included and the rest removed.

4.1.2.1 Summary

As a result of these tests, the following DNMs were removed from the dataset: *Alphabay, Andromeda, Black Services Market, Blue Sky, Cannabis Road 3, Cantina, Cryptomarket, Drugslist, Hydra, Middle Earth, Outlaw Market, the Market Place, Tochka, Tor Escrow*, and *Tor Market* and *Evolution* was partially removed.

As *Alphabay* grew to be the largest DNM in 2017, its omission from the dataset potentially obscures the recovery of the ecosystem after Operation Onymous and the Evolution Exit Scam.

For the remaining DNMs, which are *Breaking Bad, Dark Net Heroes, Dark List, Deepzon, Free Market, Grey Road, Onionshop, Pigeon, Silk Road Reloaded, the Majestic Garden, Tor Bazaar, Tortuga 2* and *White Rabbit*, no metric could be found to validate their scrapes. As such, they are also removed from the dataset. This leaves 34 scrapes that have been verified and are presumed to be a good enough representation of their DNMs that they can be used for analysis.

TABLE 4.4: Review Data Available on DNMs

Market	Type of Data	Verifier Available
Abraxas	Rating, Date, User, Comment,	None
Agora	Comment	None
Amazondark	Comment	None
Area 51	Comments	None
Bloomsfield	Comments	None
Cannabis Road 2	Comments	None
Cloud 9	Comments	None
Dark Bay	Comments	None
Diabolus	Comments	None
Dreammarket	Date, Comment, Number of Stars, Price	None
Evolution	Comment, Name of User, Date	None
Freebay	Rating, Comment	None
Haven	Rating, Comment	None
Nucleus	Rating, Comment, Price, User, Date, Product	None
Oxygen	Comment, Rating, Date	None
Panacea	Rating, Date, Vendor and Comment, Product, Price	None
Pandora	User, Rating, Comment, Date, Value	None
Pirate Market	Review, Rating, Username, Date	None
Poseidon	Rating, Review, Product, Date	The number of vendor transactions is provided but no feedback was given in 400 out 404 instances
Silkroad 2	Rating, Comment, Date	None
Silk Street	Comment	None
The Real Deal	Comment	None
Tom	Rating, Comment, Date	The DNM publishes the number of items sold for each listing but only one piece of feedback per listing
Topix 2	Date, Product, Buyer, Seller, Comment, Rating	The DNM publishes the number of reviews. There are 985 reviews and in 156 (16%) instances the number of reviews is equal to the recorded total number of reviews and in the rest the observed number of reviews is greater than the recorded total number (in 692 (70%) instances there is one observed review and none recorded).
Torbay	There are no reviews	The number of vendor reviews and number of bids for each product

4.1.2.2 Review Data

Some studies have used the reviews posted on product listing pages to approximate the number of sales taking place on DNMs (Christin (2013)). This is possible where reviews are compulsory on a DNM and therefore can be taken as a count for the number of purchases.

The previously described process of employing metadata within the DNM pages to assess the completeness of the set was applied to review data. The quality of review data was of particular concern because many DNMs only published the top reviews or limited the visible reviews by date. Further, when manually inspecting elements of the dataset it was observed that, when reviews spanned several pages (because the product had a large number of reviews or the site reserved a small space only for reviews), the scraping tool used would often not find every page of reviews, omitting a significant proportion of them.

24 of the 34 DNMs included in the dataset had review data for listings and, of these, just 4 had additional data that could be used to evaluate the completeness of the reviews. The details of the reviews are given in Table. 4.4.

The review data could not be evaluated for many DNMs and, in the few instances where it could be, the proportion of reviews found was very low. It was therefore concluded that the review data could not be used to approximate sales for this dataset.

4.1.2.3 Ecosystem Coverage

In order to estimate the proportion of listings missing from the scrape, the sum of the observed listings across all DNMs was summed per date and compared to the advertised number (for DNMs where both pieces of information was available. This produced Figure 4.2(a). The average percentage was 48.0% and the maximum and minimum were 100 and 0% respectively.

This analysis was repeated with the advertised and observed number of vendors for *Cloud 9* and *Panacea*. The proportion of vendors captured over time was plotted in Figure 4.1(c) and the average percentage was 12.3% with maximum 19.7% and minimum 0%.

To repeat this analysis for the DNMs where vendor information was available in the Grams dataset, the Schnabel estimator was proposed to approximate the total vendor population. The vendor population for each DNM in the Grams dataset and the corresponding vendor populations in the Branwen dataset were taken as the samples. The Schnabel Estimator was thought to be appropriate for use here as the scrapes were taken from the same population which is, therefore, constant where the scrapes were taken at the same time and as it was assumed that the vendor population captured by each scrape was chosen at random.

However, for the DNM *Cloud 9*, the proportion of vendors thought to be captured when comparing the observed population to the advertised population and the proportion of vendors thought to be captured when comparing that same observed population to the Schnabel estimated total population is very different. This is illustrated by Table. 4.5. This substantial discrepancy could be a result of the DNM falsely advertising a large vendor population, or it could be because the Schnabel estimator is less accurate with only two population comparisons. In case of the latter, the Schnabel estimator is not used to calculate the estimated proportion of vendor profiles captured by the scrape.

A final measure for the proportion of listings and vendors missing from the dataset is the number of corrupted files collected in each scrape. These are listing or vendor profiles that were downloaded but contained no information because the scraper was logged out, or terminated part way through collecting information, etc. Counting these provides a lower bound for the amount of missing information. Figure 4.3 displays the number of corrupted vendor and listing files collected from the DNMs across the measurement period. Clearly, less listing files were lost in the scraping process than vendor files. The number of missing files dramatically reduces at the end of 2014 which potentially implies that the data for 2015 is more reliable.

TABLE 4.5: Comparison of the Estimated Proportion of Cloud 9 Vendors Included in Scrape

Date	Comparison to Advertised Number of Vendors	Comparison to Schnabel Estimate of Vendors
9/6/2014	0.00820	0.0503
16/6/2014	0.164	0.932
22/6/2014	0.131	0.757
16/7/2014	0.160	0.932
24/7/2014	0.169	0.934
28/7/2014	0.156	0.946
8/9/2014	0.165	0.942
21/9/2014	0.162	0.938
26/9/2014	0.159	0.910
29/9/2014	0.164	0.926
17/10/2014	0.157	0.793
1/11/2014	0.165	0.865

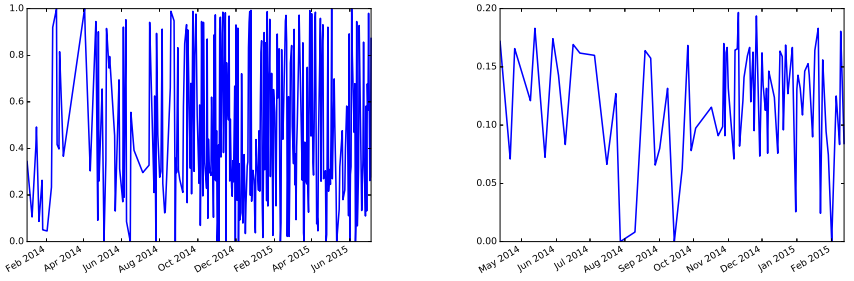

(a) **Estimated Proportion of Listings Captured** (b) **Estimated Proportion of Vendors Captured**

FIGURE 4.2: **Estimated Proportion of Vendors and Listings** Captured for All DNMs in Dataset.

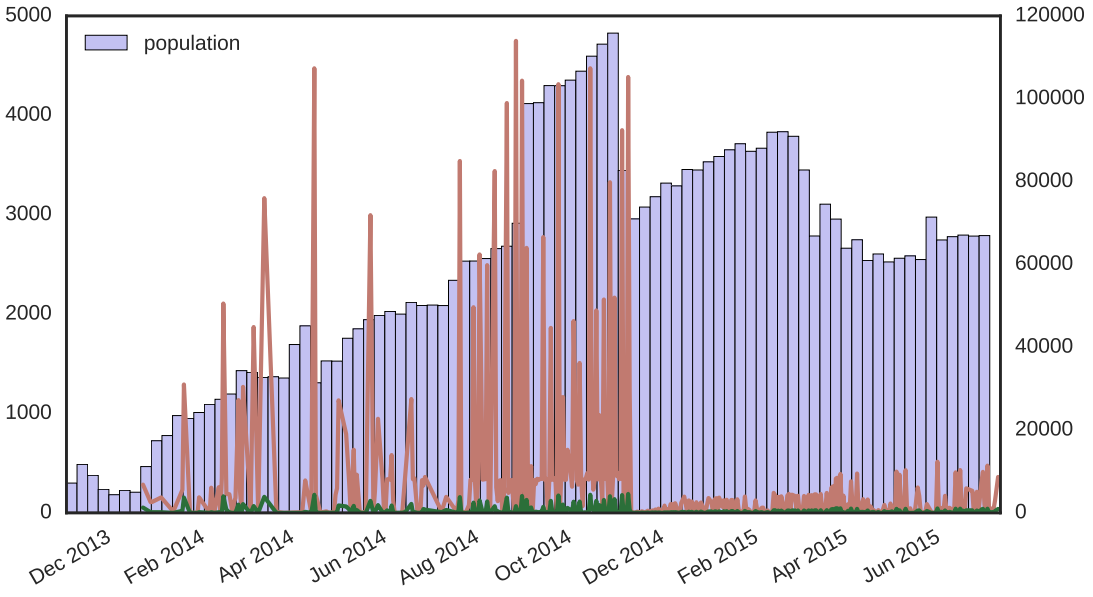

FIGURE 4.3: **Number of missing files** plot showing the number of missing listing files (green) and vendor files (pink) against the vendor population (purple).

To assess the proportion of the ecosystem captured in the study, the number of DNMs was compared to the estimated number of DNMs active over the measurement period. Graphical representation of this is presented in Figure 4.4. The number of active DNMs not present in the study is an estimate calculated using a number of sources including the independent researcher Gwern Branwen and the DNM monitoring site `www.deepdotweb.com`.

Figure 4.4 clearly shows that the estimated number of DNMs fluctuates over time in a similar manner to the number of DNMs in the dataset. This is confirmed by the Kolmogorov–Smirnov test which produced a test statistic of 0.0465 and p–value of 1.00. Therefore, the size of the ecosystem, measured by the number of DNMs, in the dataset is likely representative of the growth of the total ecosystem.

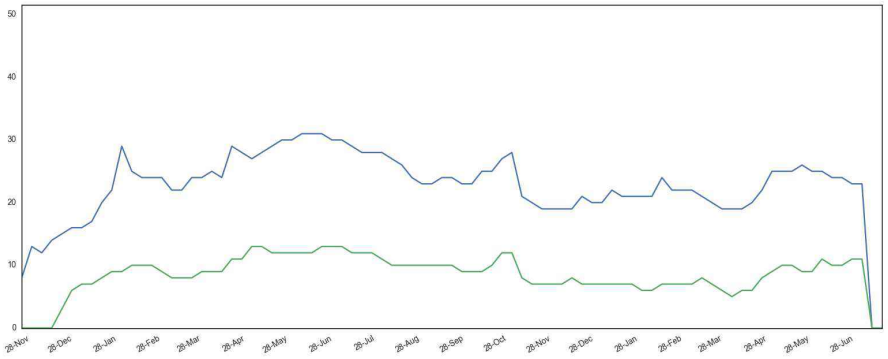

FIGURE 4.4: **Estimated Ecosystem Caputer** expected number of DNMs (blue) vs Observed Number of DNMS (green).

Because the size of the DNMs not in the dataset are not known, the size of the ecosystem measured by the size and popularity of the DNMs cannot be calculated.

4.1.3 Forums

Forums were collected on the following DNMs: *Abraxas, Agora, Andromeda, Black Market Reloaded, BlackBank Market, Bungee54, Cannabis Road 2, Cannabis Road 3, DarkBay, Darknet Heroes, Diabolus, Doge Road, Evolution, Gobotal, GreyRoad, Havana/Absolem, Hydra, Kingdom, Kiss, Mr Nice Guy 1, Nucleus, Outlaw Market, Panacea, Pandora, Pigeon, Project Black Flag, Revolver, Silk Road 1, Silk Road 2.0, TOM, The Cave, The Hub, The Majestic Garden, The RealDeal, TorEscrow, TorBazaar, Tortuga 1, Underground Market, Unitech, Utopia*

The forum scrapes contain threads on a number of different topics including drug quality, safe usage practice, and vendor reputation. Different scrapes have stored data in different formats but, broadly, each contains the threads accessible at the time of the scrape, the date they were started, the user who started the thread, the users who commented on the thread, their comments, and the date and time they were posted.

In addition to the forums maintained for each DNM, users also used public sites such as `Reddit.com` to discuss DNMs (Reddit User (2017)). Data from Reddit was also collected and is outlined in Section. 4.3.

4.2 Silk Road 2.0 Dataset

This dataset was collected from the servers of the DNM *Silk Road 2.0* by a UK LEA. The DNM was operational between 4 November 2013 and 6 November 2014. It contains information on 24,242 accounts identified as belonging to UK buyers and 151 accounts identified as belonging to UK vendors. For each user, the following data is stored in .xml files:

- **User Data** the username, type, and profile description of the user is recorded as well as the date they first and last logged in, the number of times they logged in, the country they claim to reside in, and the countries they ship to (if they are a vendor).
- **Transaction Data** each transaction is recorded with the product name, price, vendor, buyer, date, and transaction status (i.e. has the transaction been completed);
- **Messaging Data** each message that the user sent or received is recorded with the sender, receiver, subject, date, and message contents;
- **Deposit and Withdrawal Data** the amount the user deposited and withdrew from their account is recorded with the date of the withdrawal/deposit and the wallet ID to which the coin was moved/ from where the coin was transferred.

In addition to these complete profiles, information on users from outside of the UK could be constructed from the information in the dataset. 767 additional user accounts were observed selling products and a further 17,037 accounts were observed buying products. For each of these users, estimates of their lifetime and transaction history could be made, but for their activity involving UK based buyers or vendors only. The approach for making these estimates is explained in Section. 5.

The dataset contained information on 190,802 transactions, of which, 14,874 were declared “cancelled”, 136,522 were declared “shipped”, 1,025 were declared “ordered”,

and 38,381 "finalized". 96,558 of the transactions were between buyers and vendors with profiles in the dataset, the remaining 94,244 transactions were between 903 users with profiles in the dataset and 17,804 users who did not have profiles in the dataset (because they were presumed to be operating from outside of the UK).

The data is incomplete as some messages have been deleted (either when the account they were sent from was deleted or when the user chose to delete them themselves) and so only retain sender, receiver, time, and subject heading data.

This data contains some private information such as names and addresses, as well as BTC addresses which could potentially be linked to real world identities. It is stored in the secure UCL Data Lab and ethical approval for conducting research on this data has been granted. Information on users was anonymised by replacing usernames with unique number identifiers and information that could not be anonymised, e.g. messages, was not removed from the lab.

4.3 Reddit Data

The subreddit */r/darknetmarkets* was created in December 2015. It was banned in March 2018 for contravening Reddit's rules on facilitating the sales of illegal goods and services. The forum was used to discuss topics related to Dark Net Markets and so provides a useful resource for researchers to understand how users felt about events, both minor and major, and how they discussed them on the Clear Web.

The posts and comments made on the subreddit */r/darknetmarkets*, alongside the other contents of Reddit, were collected by Reddit user Stuck_In_the_Matrix and made available for research in 2015 (Stuck_In_the_Matrix (2015)). The posts and comments can be accessed through BigQuery, an online data analysis facilitator[2]. The repository is separated into comments and posts with a dataset for each month. From this repository, the posts and comments made on */r/darknetmarkets* between 1 September 2016 and 30 November 2017 were extracted and combined to rebuild the subreddit. This process returned 40,353 posts and 581,465 comments.

For each post, the time the post was created (*created_utc*), the author, the number of comments received (*num_comments*), the score and the number of up and down votes the post received, the contents of the post (*selftext*) and the name of the post were collected. For each comment, the comment text (*body*), author, time created, reference code for the post the comment was made on (*parent_id*), number of up and down votes and controversiality score received, and name of the comment were collected.

[2] https://cloud.google.com/bigquery/

The data was aggregated in two phases, first the comments were attributed to the relevant posts in each monthly dataset and secondly the monthly datasets were combined. Comments were attributed to posts by matching the comment *parent_id* and post *name*. The monthly datasets were combined to create a dataset with an entry for each post which contained all comments across its lifetime, the date it first appeared in the dataset and the date it was last amended, the most up to date score and number of up and down votes in the repository, the contents of the post and its author.

This process of aggregation increased the number of posts from 40,353 to 324,120 and reduced the number of comments from 581,465 to 572,585. 313,225 (97%) *parent_id*'s were missing from the post dataset and so, for these comments, artificial posts with no content were created. The data on the number of comments per post was used to estimate the number of comments missing from the dataset. This metric could only be used for posts that were included in the dataset. 4,088 (38%) of the posts found in the dataset were either of the form '[deleted]' or '[removed]'. The predicted number of comments in the dataset was 137,097 and the recorded number was 53,196 (39%).

These measures of completeness are underestimates as, through inspection of the content, comments that clearly linked to other posts (by referring directly to unique scenarios) were identified. This implies that some comments are replies to comments on threads but that either the threads they are replies to are missing or that an *id* variable was not collected in the dataset.

The number of contributors to the subreddit was calculated by counting the number of usernames recorded in the dataset. Where the author of a post was not known (i.e. in the instances where a comment was found such that there was no corresponding post with matching *parent_id*) the author was given a new id corresponding to the id of the post. This approach prevents the conflation of posts by anonymous authors but overestimates the number of contributors.

There were 347,844 unique contributors names, however one such name was '[deleted]' which is presumably not a user name but instead a reference to the real identifier being removed from the site. 69,858 contributors had the name '[deleted]' and so they were each given a unique id. As before, this prevents posts by different contributors being conflated but may lead to an overestimate of the number of contributors. It is more likely that the posts/comments are not connected given that the next highest poster had made 139 posts and had the name 'AutoModerator' and so is presumably an automated response from the site to help manage subreddits. Separating each deleted account gave a total of 417,701 contributors.

In addition to the subreddit */r/darknetmarkets* the subreddit */r/dnmuk*, which was banned at the start of 2018, was collected to give a UK specific perspective.[3] Initially, this subreddit contained 18,890 posts and 285,791 comments. These were combined to create a dataset of 168,873 posts and 281,248 comments. 166,830 (99%) of the posts were artificially created from comments and, of an expected 21,997 comments, only 8,602 (39%) were found. 549 (26%) of the posts found were either of the form '[deleted]' or '[removed]'. There were 173,150 contributors found in the dataset, of which the username '[deleted]' was observed 27,643 times. Separating this username into individual accounts increased the estimated number of contributors to 200,792.

In order to compare the behaviour on these subreddits to a similar subreddit which does not focus on illegal activity, the subreddit */r/Ebay* was collected for the same time period. According to Reddit, this subreddit has 26K users and has been active for 10 years (Reddit (2018)). The dataset created from this subreddit initially contained 20,558 posts and 88,076 comments. It was aggregated into a dataset with 54,606 posts and 87,264 comments. 759 (16%) of the posts were either of the form '[deleted]' or '[removed]'. 49,846 (91%) posts were created from comments and, of an expected 18,770 comments, only 8,676 (46%) were found. There were an estimated 58,630 contributors in the dataset, 5,798 of which had the username '[deleted]', separating this username increased the number of contributors to 64,428.

[3]Multiple other subreddits exist (and have existed) to facilitate DNM themed discussions. It was found that the analysis on these two threads were exhaustive and so the data was not supplemented with more subreddits.

Chapter 5

Methodology

This section describes the methodology used to answer the hypotheses outlined in the previous chapter. It details three studies: a cross network study that quantitatively evaluates the impact of Operation Onymous, a study of buyer behaviour on the DNM *Silk Road 2.0*, and, a qualitative study of Reddit discussions about Operations Hyperion and Bayonet and the closure of Hansa.

5.1 Cross Market Study

The raw data from the Gwern Brawnwen Dataset was transformed into a database containing 10,092 vendors. Each entry corresponds to a vendor's profile on a DNM and contains their username(s), PGP key(s), the date they registered, the date they were last seen, the last date their profile was active, the products they had listed, their reputation, and their profile description. As this information was presented differently on different DNMs, some steps had to be taken to homogenise it.

The Cross Market Study was used to evaluate the impact of Operation Onymous. As such, vendor behaviour, specifically whether vendors continued to trade after the incident and, if so, which DNM they transferred their business to, was measured.

Each of these processes is described below.

5.1.1 Linking Vendor Accounts Across DNMs

Some vendors manage profiles on multiple DNMs (Soska and Christin (2015); Broséus et al. (2016)). This may be a strategy designed to minimise the cost of market disruption (Soska and Christin (2015)) or because operating on different DNMs allows for a greater

diversification of products (Broséus et al. (2016)). In either case, it is necessary to be able to link accounts belonging to the same vendors in order to more accurately estimate the overall population, as well as better understand the economic strategy of vendors.

The first step of this analysis was conducted on all of the DNMs, but the linking of vendors through their profile descriptions and products was only conducted on the validated datasets.

Four features were identified as being useful for linking accounts: usernames, PGP keys, profile descriptions and product listings (Broséus et al. (2016); Soska and Christin (2015); Iofciu et al. (2011)). Two accounts on different DNMs were considered to be owned by the same vendor if they shared one of the following:

- a **username** exclusive of case (i.e. lower case and upper case were considered equivalent), spaces or punctuation characters (e.g. '-' or '_');
- a **PGP Key** but only in the instance that their usernames are similar, i.e. their usernames have a Levenshtein Ratio of 0.75*[1] excluding cases where the Longest Common Substring (LCS) is a common word.
- a **LCS** such that the substring was equivalent to one of the usernames exclusive of case, spaces, punctuation characters or common words.

PGP Keys could not be used to exclusively match vendors as some vendors listed a site key on their profile, instead of a unique key, leading to many vendors on the same site being mistakenly linked. It was, therefore, necessary to also compare usernames using the Levenshtein Ratio. This metric of similarity identifies vendors who have had to adapt their username on a different DNM but have made it recognisable, for example "DrWhite" and "DrWhiteTeam", by matching usernames that contain similar patterns. The Levenshtein Ratio counts the number of changes required to turn one string into another and presents this proportional to the size of the string. Usernames that are only a few characters different relative to their length are considered similar.

Matching vendors using this method can lead to incorrectly matching vendors who use common words, such as "cannabis" or "drugs" in their username. To account for this a set of common words were created and usernames were not matched if their Longest Common Substring (LCS) was a common word. The set of common words was created by extracting the most frequent LCS of 5 characters or more across all of the usernames. Of these, the top 663 were chosen, the full list is in Appendix D. Each of these strings was used in at least 53 usernames. This set of variables was found experimentally by testing the method on a ground truth dataset.

[1]The threshold of 0.75 was chosen experimentally by varying the threshold between 0.5 and 0.95 in increments of 0.05 until all matches in the ground truth dataset were identified.

The ground truth dataset was created using vendor profiles on the DNM Darklist. This DNM allowed vendors to list verified accounts on other DNMs on their profiles. 50 vendors stated they were also operating on *Tormarket* and so the above method was applied to the vendors on these two DNMs. Matches not in the ground truth were manually inspected and their profile descriptions and listings were compared. If two profiles listed at least some of the same products, explicitly referred to each other's account or shared sentences or phrases in their profile descriptions, the accounts were considered to belong to the same vendor. If none of these features were present, the accounts were considered to be distinct. As a result of manual inspection, 17 matches were added to the ground truth. The size of the set of common words was varied until the method found all the matches, with the minimum number of false positives, which was two. The method incorrectly matched the accounts "fake" and "fakemarket" and "rad" and "colorado".

After this matching process, all vendors operating on more than one site were manually inspected, of which there were 1,225 linked to 16,692 usernames. In each instance where the usernames were not a precise match, the profile descriptions and product listings were compared as before. 208 vendors with 14,435 alternative usernames were found to be false positives and so were removed from the database before their information was collected.

These characteristics are ones that can be identified in accounts vendors have designed to be easily connected. However, some vendors may create accounts that deliberately look different to other accounts they own, for example because they have a poor reputation on a site and want to start again. Even if a vendor creates a new account on a different DNM that they do not want to be linked to their existing accounts, some characteristics may remain consistent across all of their accounts. These characteristics are the way that they write their profile description and the products that they sell (Iofciu et al. (2011)).

To compare vendors using their profile descriptions, the cosine similarity was used to measure the distance between vector representations of the Term Frequency - Inverse Document Frequency (TF-IDF) analysis of each description. TF-IDF vectorisation extracts words from a text and assigns each a value based on how frequently they appear in the text weighted by how frequently they appear in the entire corpus, i.e. a word that is frequently used in one text but is used scarcely in other texts in the collection will score more highly than a word that is commonly used in every text.

The formula for the TF-IDF value of a word v in a text t from a corpus C is as follows:

$$\text{TF-IDF}(w) = \frac{n_w}{|t|} \cdot \log(\frac{|C|}{N_w})$$

s.t. n_w is the number of times that the word w appears in the text t and N_w is the number of texts that contain the word w in the corpus C (tfidf.com (2018)).

This analysis was conducted using a prebuilt pipeline from Sci-kit Learn[2]. The cosine similarity measures the distance between two vectors by calculating the cosine of the angle between them. The cosine similarity of a vector compared to itself is 1.

Vendor accounts that did not have a profile descriptions were excluded. As were any accounts that had descriptions less than 100 characters as they tended to be a variation of "this vendor has no profile description".

First, for each vendor active on multiple DNMs, the descriptions associated with their different profiles were compared pairwise. The pairwise cosine similarities were then combined as a mean value for each vendor. The average of this value per vendor was 0.680. This value was used to identify potential matches in the rest of the vendor population by comparing the TF-IDF of each profile description to that of every other and recording the matches greater than 0.680.

385 vendor matches were found from this process. Of these, on 160 of these matches, the Levenshtein ratio between the two usernames was 0.75 or greater and, for 52, the string of one username was fully contained in the string of the other. This implied that the process had found matches between profiles which had username based indicators that were not captured by the other matching mechanisms. For example, matching *justsmuggledn420* and *JustSmuggledN*.

As relying on similarities in usernames will fail to identify accounts matched created to look different, one final measure was applied to those matched based on their products. Whilst it is not necessarily the case that vendors sell the same products on all DNMs (Broséus et al. (2016)), if two vendors sell the same list of products on more than one site, this can suggest they are run by the same people offline. To match vendors based on their products, all of the products that a single vendor sold were identified and then compared to the set of products sold by every other vendor from the ecosystem.

The products sold by each vendor were compared by calculating the number of products sold by both vendors as a proportion of the total number of products they sold between them. If this value was greater than 0.224 then the two vendors were inspected further. To find this lower bar, the vendor accounts that had been previously matched were

[2] https://scikit-learn.org/stable/

compared using their products. For each vendor, the products sold through each separate account they owned were collected and compared pairwise. The largest proportion value was taken for each vendor and the average of all these values was used to find the value 0.224.

This condition was not sufficient to confirm matches alone as some vendors gave products generic or commonly used names. For example, if two vendor accounts are selling *Cocaine* this does not necessarily imply they are controlled by the same vendor. To account for this, if a vendor account only sold one product, that product was manually inspected. If the product name only contained brand or item names and weights/sizes (as opposed to, for example, the vendor's username or a unique punctuation combination) then they were not matched with any other accounts.

Some vendors sold hundreds of different products and so their inventories could not, realistically, be manually inspected and evaluated based on the genericness of their products. Instead, the usernames of potentially matches were compared. If two accounts had similar usernames, e.g. they were the same word or phrase using a different spelling, or the suffix of one was a pseudonym of the suffix of the other, etc., then the accounts were considered a match (see Appendix B for all of the matched usernames).

This reduced the size of the database from 10,152 to 10,092.

There is no ground truth[3] that can be used to evaluate this methodology and so it cannot be concluded that all of the vendor groups were found. Therefore, the size of the ecosystem reported is likely an underestimate. However, it can be concluded that the method of linking vendors solely by their usernames and PGP keys missed at least 429 vendor pairs and therefore had a false negative rate of at least 4% for this dataset.

5.1.2 Measuring Vendor Lifetimes

Where possible, the lifetimes of vendors were measured using information provided by the DNMs. However, whilst some DNMs provide the date for when a vendor created their profile (Registration Date) and the last date that they were on the site (Last Seen Date), this was not so of all DNMs. In these instances, these dates were approximated as the first and last dates the vendor profile was present in the dataset, respectively. The precision of the Last Seen Date could not be verified - it could not be evaluated whether or not this date referred to the last time the vendor logged on and if it was

[3]It was noted in the viva that a potential ground truth is available in the form of a database created from arrest records. This database has the potential to link seemingly unrelated accounts, e.g. HumboldtFarms and PureFireMeds to the same vendor team. However, of the arrest records that were sourced in the course of this project, none could be connected to accounts included in the database.

updated if the vendor stayed logged in for several days or if it was the last date the vendor carried out an action on the site (e.g. made a sale, deposited funds, edited their profile, etc.) and if this was consistent on different DNMs. As such, in addition to the collection of this date where available, the last date the profile was present in the dataset was also collected. For 87% of vendor profiles, this date did not differ for the different measurements. Further, using these two measures of lifetime did not produce a difference in the analysis and so only the former measure for the last date was used.

5.1.3 Collecting Product Information

Examining the product listings that each vendor advertises can provide some information about their economic strategy and investment in the ecosystem. Vendors who have no products for sale may have created an account to explore the DNM but decided the risk was too high to begin trading. Vendors who consistently sell the same product type during their lifetime may have a consistent supply whereas vendors who constantly change their inventory may have a more opportunistic method for acquiring stock. Some vendors sell different products on different DNMs (Broséus et al. (2016)), perhaps implying a deliberate market strategy. The pricing information of each product is also useful as it has been shown that vendors are able to charge a premium if they have a higher reputation and so product prices can reflect a higher quality product and/or service (Hardy and Norgaard (2016)). Some DNMs allowed vendors to advertise their listings directly on their profiles, though this was not true of all. Where this was not the case, products were collected from their listing pages and linked to vendors using their usernames.

Whilst all trading occurred in cryptocurrencies, the products were advertised in a number of different currencies including U.S. Dollars, Canadian Dollars, Pound Sterling, Bitcoin, Litecoin, and Dogecoin. All prices were converted into U.S. Dollars and the exchange rates for the different currencies during the time period were found on `investing.com`, `coindesk.com`, `crytocoincharts.info`, and `coinmarketcap.com` (Fusion Media Limited (2016a,b,c); Limited (2016); CoinMarketCap (2016); Coin.Market (2016b,a)).

DNMs offer a wide variety of products (Christin (2013)) and vendors are able to name and describe their products themselves. As such, products were first grouped into categories before they were compared. The categories were as follows:

- **Benzos** - a family of tranquillisers containing Benzodiazepine, including prescription Benzodiazepines.

- **Cannabis** - any product made from or with parts of the cannabis plant (including Edibles, Wax, and Oil) and any synthetic product created to reproduce the effects of cannabis, i.e. Synthetic Cannabinoids.
- **Dissociatives** - any hallucinogenic which produces a feeling of dissociation and is not a cannabis product, Benzodiazepine, Opioid, or Barbituate.
- **Ecstasy** - any MDMA or MDMA based product, including MDA and ethylone based products.
- **Opioids** - any opium derived product or synthetic product designed to reproduce the effects of Opioids, includes Prescription Opioids.
- **Prescription** - any product that requires a medical prescription not including Prescription Opioids, Stimulants, Steroids, or Benzodiazepines.
- **Psychedelics** - any product that produces hallucinations excluding Ecstasy, Benzodiazepines, Dissociatives with hallucinogenic effects or Barbituates.
- **Steroids** - any artificially produced hormones including Prescription Steroids.
- **Stimulants** - any Cathinone based product or 'Designer Drugs' that produce stimulant effects including Prescription Stimulants.
- **Other Drugs** - any drug based product not already categorised including Alcohol, Barbituates, Research Chemicals, Tobacco, Weight Loss products, and Custom Listings.
- **Digital Goods** - any product that exists only in digital form, for example eBooks or hacking software.
- **Other** - all other products.

These categories were created by comparing all of the different categories across all DNMs and finding the most utilised (Lee and Antin (2011)). Any drug category featured on the majority of DNMs in the study was included in the set of categories, all the others were categorised as "Other Drugs". To establish the boundaries of each category and remove contradictions, the subcategories from each DNM were sorted into the category set. Where subcategories had been sorted into multiple categories, the majority definition was taken. A full list of the categories and their subcategories is listed in Appendix C. Because these categories are different to those used in other, similar work, a new classifier had to be created that sorted the listings into the categories.

To sort the products into the different categories, a Stochastic Gradient Descent Classifier was trained on TF-IDF features extracted from the listing titles (Pedregosa et al. (2011)). The analysis was conducted using Python 2.7. A ground truth dataset of 3,000 products (50 products randomly chosen from 60 DNMs) was hand coded[4]. 90%

[4]The products were only coded by one coder and so reliability could not be measured through inter-coder agreement. However, the products were coded by verifying what kind of drug they were using several drug databases and so the coding practice was not open to interpretation.

TABLE 5.1: Category Classifier Accuracy

	Precision	Recall	F1-score	Support
Psychedelics	0.88	1.00	0.94	22
Steroids	0.83	0.71	0.77	7
Other Drugs	0.89	1.00	0.94	17
Ecstasy	1.00	0.75	0.86	4
Dissociatives	1.00	0.91	0.95	22
Digital Goods	0.92	0.96	0.94	80
Stimulants	0.89	0.94	0.92	18
Prescription	1.00	0.75	0.86	4
Opioids	0.97	0.85	0.90	39
Cannabis	1.00	0.44	0.62	9
Benzos	0.91	0.59	0.71	17
Other	0.82	0.97	0.89	60
Avg / total	0.91	0.90	0.89	299

of the ground truth data was used for training and the remaining 10% for testing. The products were handcoded using the resources `drugs.com` and `talktofrank.com`. The classifier reported an overall F1 score 0.89, for more precise results refer to Table. 5.1.

Some categories, such as Benzos, Opioids, and Stimulants contain Prescription Drugs, and so the classifier was less performant in these four categories. There are thousands of strains of cannabis, some of which are named after other types of drug, for instance LSD (Leafly (2017)). Similarly, there are many types of Ecstasy, some of which do not refer specifically to Ecstasy or MDMA in their title, e.g. Nintendo Pills (Enlighten (2017)). As such, the classifier was more likely to mislabel these products.

5.1.4 Measuring Vendor Reputation

An attractive feature of DNMs is the ability to provide feedback on vendors and read the feedback of others (Van Hout and Bingham (2014, 2013b)). Many DNMs provided some metric to amalgamate vendor feedback into a reputation score to more easily compare vendors. This reputation value, and how it changes over time, can be an indicator of success of a vendor. It may also be a determining factor for a vendor to leave a DNM - if their reputation falls to a sufficiently low level, it may be worthwhile deleting their account and starting again, alternatively, a high reputation may make moving to a different DNM a worse financial decision if they are unable to transfer their reputation.

Because different DNMs have different ways of measuring and displaying reputation a metric was needed that could describe the status of a vendor on one DNM in a way that was both comparable to other vendors on the same market and to the same vendor on

different markets. Instead of comparing the absolute reputation variable provided by the DNM, this information was transformed into a vendor ranking which described how good that vendor's reputation is in comparison to all the other vendors on the DNM. This measure was calculated by comparing the reputation value of every vendor on a DNM and producing an ordered list such that the vendor with the highest reputation value is ranked as number 1.

On some DNMs the reputation values exhibited little variance and many vendors had, for example, a reputation of 100% because they were new vendors. To encode this information into the rankings, a variable was added to each ranking equal to the number of vendors who shared that ranking divided by the number of vendors on the market at that time i.e.

$$R_v = r_v + \frac{n_r}{V} \tag{5.1}$$

s.t. R_v is the rank value of a vendor, r_v is their ranking when all of the vendors are ordered, n_r is the number of other vendors who share this ranking on their DNM, and V is the total vendor population on the DNM on that date. This measure ensures that a vendor who is the only vendor in their rank position on their DNM will score higher than a vendor who shares their rank position with several other vendors on their DNM.

A number of different reputation measures were used to produce rankings. Some DNMs provided a value for each vendor and so this was used, where available. When an overall value was not available, other information such as the number of positive reviews or successful sales as a proportion of the total number of sales was used. Whilst there are multiple ways of combining the sales information, this one contains information on positive sales and total number of sales and was common rubric used by other DNMs. In instances where multiple pieces of information were available, the self defined reputation value was preferable as it represented the information considered to be important by the DNM itself. If this information was not available, the metric which provided the greatest variability was used. When only review data was available, the reviews were transformed into numerical values using a sentiment classifier. This method was the least preferable as, to create the classifier, subjective decisions had to be made about whether a review was positive or negative. When a DNM presented multiple possible values but none were available for every vendor, the rankings were calculated from the reputation value used by the most vendors. For more specific detail on the reputation calculation of each DNM, refer to Appendix E.

To determine whether reviews were positive or negative, 500 reviews were hand coded as either positive, neutral or negative by two individuals. A review was positive if it contained a positive comment about the product or vendor, neutral if there was no positive or negative statement present, and negative otherwise. This corpus was

used to train an SDG Classifier from the Python 2.7 library sklearn. To mitigate the low proportion of negative reviews in the training set, the initial classifier was used to identify a further 251 negative reviews which were then added to the set as well as 100 more positive reviews, increasing the set to 851 reviews. The reviews were prepared through TF-IDF analysis. The classifier had an F1 score of 0.89 and 91% recall. For every review a vendor had, 1 point was added to their reputation if it was positive, 0 if it was neutral and -1 if it was negative. If reputation information was not available on a given date, the vendor's reputation was taken to be their most recent reputation value.

5.1.5 Identifying Events

A timeline of events was collated that might have affected the ecosystem between April, 2010 and December, 2016. The time line includes 65 DNM openings and 61 closures, 3 attacks on Tor and 18 attacks on DNMs (such as hacks or Doxxings), 15 arrests, 7 drug related law changes, 20 media events (such as articles about DNMs in mainstream publications) and 13 miscellaneous events. It was created using information collected by independent researcher Gwern Branwen (Branwen (2013)) and by monitoring relevant news websites such as `deepdotweb.com`, `reddit.com`, `wired.com`, and `forbes.com`. The full timeline is given in Appendix F.

When measuring the impact of events in the time line, the vendors affected are considered to be those active on the ecosystem in the week before the event started and during the event itself. The easiest way to identify the start and end dates of an event are to look at the recorded dates. The start date of Operation Onymous was identified from reports as on 5 or 6 November 2014. As such, 4 November was chosen to ensure the start of the event, regardless of time zone. The end date was chosen to be one week after this date, i.e. 11 November 2014, because this time period included every DNM shut down during Operation Onymous as well as the last date of their scrapes. Therefore, the vendors affected by Operation Onymous are all those active from 28 October until 11 November 2014.

The vendors affected by the Evolution Exit Scam are those operating between 28 February and 18 March 2015. The exit scam took place between 14 and 17 March 2015 however, prior to this, many users were unable to withdraw coin from their account (DeepDotWeb (2015d)). As such, the event can be considered as taking place when users first begin leaving and not when Evolution closed. This was identified by inspection as being 7 March 2015.

5.1.6 Continuing to Trade

When an account or entire DNM is shut, the vendors who are present during the event have two possible actions. They can cease trade and leave the ecosystem or they can continue to trade. In order to continue to trade, they can maintain accounts that are operating on unaffected DNMs and/or they can create new accounts on unaffected DNMs. The new and maintained accounts were identified by linking the usernames of vendor accounts, as described above.

In addition to determining how many vendors continue to trade after an event, the time taken by vendors to continue to trade can also be calculated. This is taken to be the time taken for a vendor to create their first new account after their account(s) is shut during an event, i.e. the period for which they are inactive after an event.

5.1.7 Vendor Movement

Exponential Random Graph Models (ERGM) are tools of network analysis that treat the edges between nodes as random variables (Robins et al. (2007)). If it is considered that the network being studied is one possible configuration out of many different possible ways of connecting its nodes then probabilities can be assigned to each of the edges that correspond to likelihood they would be formed in the other variations of the network. ERGM is a method of finding those probabilities and, as such, can be used to understand why the network structure has formed in a particular configuration (Robins et al. (2007)).

For example, an ERGM can be used to determine the increased probability that two nodes will be connected, given that they share an adjacent node. If this probability is high, that means that this closing of triangles is a significant driver of connections within the network.

To measure vendor movement after Operation Onymous and determine if any characteristics of the surviving DNMs particularly attracted users, the ecosystem was modelled as a network. Each DNM was represented as a node connected by edges weighted with the number of users who moved from that DNM to another. ERGM analysis was used to see if the size or age of the DNMs encouraged more users to move to them.

5.2 Silk Road 2.0 Study

In order to analyse the social dynamics of the DNM *Silk Road 2.0* and investigate how users select vendors and respond to different vendor behaviours, a network was

constructed from the data. The following subsections outline the process use to create the network, collect the relevant variables on each user, and conduct the various network based analyses.

5.2.1 Creating the Network

A network was created with nodes representing the 42,197 users (24,242 UK buyers, 151 UK vendors, 767 external vendors and 17,037 external buyers). Transactions were represented by directed edges that were drawn from the buyer of the item purchased to the vendor, i.e. a_{ij} represents the purchase made by i from j. The network was built using the iGraph library for Python 2.7 [5].

In addition to looking at the network of all users and all buyers, temporal snapshots were taken at weekly intervals. These networks featured all of the users whose account was active during the week, i.e. all users who had an account either starting in or before or ending in or after a given week and all transactions that took place during that week.

5.2.2 User Information

For the users whose profiles were included in the dataset, extracting information was simple as it was already organised in .xml files. For each user, it was known which date they first created their account and the last date on which they logged in, these were used to calculate the user's lifetime. There was also a list of all of their purchases which included the item they purchased, the date they purchased it on, the vendor they purchased from, and their review. For most transactions, a numerical score (out of 5) was also assigned to the purchase, as a rating, however this was not the case for all purchases.

In order to assign ratings to purchases that did not have ratings, a classifier was built using Python 2.7 NLTKs pre-built Nave Bayes Classifier [6]. The classifier was trained on whole reviews that had ratings and reported an accuracy of 94.8, an accuracy deemed sufficient that the classifier could be used without further investigation of features. All 48,060 unrated transactions were assigned a rating of 5. This was unsurprising as 93% of the transactions that did have ratings were rated 5/5.

The location for each of these users was assumed to be the UK, however the methodology for how UK users were identified by the law enforcement agency who provided the data is unknown.

[5] http://igraph.org/python/
[6] http://www.nltk.org/_modules/nltk/classify/naivebayes.html

Users were assigned the label of buyer if they made at least one purchase, even when this disagreed with the labels included in the dataset. Users were assigned the label of vendor if they sold at least one item, again, even when this disagreed with the labels included in the dataset. This is because the method used to assign these labels is unknown but suspected to be the kind of account opened by the user. This is not known for all the users in the dataset (i.e. the ones from outside the UK who do not have profiles) and so these definitions were used for consistency in the analysis.

It was possible for users to be a vendor and a buyer, or neither. In total there were 33,123 buyers, 16,086 of which had profiles in the dataset and, of these, 15,981 were labelled as buyers in the dataset and 105 as vendors. There were 909 vendors, 142 of which had profiles in the dataset and, of these, 139 were labelled as vendors and 3 as buyers.

For users who were not in the dataset, the relevant information had to be constructed from transaction data. These users did not have profiles as they were assumed to be from outside the UK but could be identified from their purchases with UK users. Each transaction that these users participated in was collected, the date of the first transaction was taken as an approximate start date and the date of their last transaction was taken as an approximate end date for their lifetime. These values are estimates and so, when conducting analysis on the lifetimes or number of transactions of users, only the users with profiles in the dataset were used.

To determine a location for each of these users, the shipping option for each transaction was examined. For instance, if the shipping option specified a location (e.g. 'Free 1st Class Royal Mail Europe') that was taken as the location of the buyer. Shipping options that included the word from (e.g. 'From Germany') were used to determine the vendor locations. Where the words 'inside' (e.g. 'inside EU') or 'domestic' (e.g. 'Free domestic shipping within Norway') were used, it was presumed that both the buyer and vendor were located in that country. Finally, where the word 'outside' was used (e.g. 'WW Airmail (outside UK)') it was presumed that the vendor was located in the country stipulated.

Assumptions were not made about the postage method used, e.g. it was not presumed that a Royal Mail or First Class option referred to the UK postage system as this could not be validated without a ground truth dataset. It was, however, presumed that users selected the correct shipping option for their item, e.g. if they were not based in the UK, they did not select domestic UK postage.

It also cannot be accounted for when one user makes multiple transactions in different locations, e.g. because they make purchases whilst on a trip. As some users made

multiple transactions that presented conflicting information on their location. In these instances, the data was aggregated to a larger geographical area in order to resolve the conflict, e.g. if someone made a purchase that implied they were in France and then a second purchase implying they were in Germany, they would be labelled as being based in Europe. The rules for this aggregation were as follows:

- Australia and New Zealand → Oceania
- Ireland and UK → Uk/Ireland
- Canada and the USA → North America
- Europe and any country in Europe → the specified country
- 'World Wide' and any specific country → the specified country
- Multiple countries across more than one continent → 'Unknown'
- 'World Wide' → 'Unknown'

107,134 (56%) of the transactions captured in the network had location data. 106,239 transactions had location data pertaining to the location of the buyer and 9,560 had data pertaining to the location of the vendor. 47 locations were identified, however not all locations were mutually exclusive, for example some shipping options were worldwide, some shipping options referred to multiple locations, and others were defined by the countries that were not available, e.g. the product could not be shipped to Australia and so the buyer is presumed to not be based in Australia.

To evaluate this approach, it was applied to the shipping option selected for transactions conducted by known UK users (i.e. those with profiles in the dataset). Location data could be extracted for 12,150 (50%) of these users. 10,775 (89%) of these users were identified as having picked a UK shipping option or shipping an item from the UK. 61 (0.5%) of users had transactions with a shipping option that implied they were not based in the UK and the remaining users had transactions with shipping options that implied they could be based in the UK but weren't necessarily (e.g. they had an item shipped to Europe or a country that was not Australia).

Of the 17,776 users who did not have profiles in the dataset (and were therefore presumed to be non UK based) 2,984 were assigned UK as a location and 163 were assigned UK and Ireland as a location and so this methodology has potentially misidentified the location of between 17% and 18% of users. Alternatively, the method that was used to identify UK users by the original owners of the data was flawed but this method is not known and so cannot be commented upon. A final possibility is that these users made purchases in the UK but are not UK based, for example because they made purchases when on holiday.

TABLE 5.2: Location of Other Users

Location	Number of Users		
	All	Buyers	Vendors
Asia	0 **(0)**	0 **(0)**	0 **(1)**
China	1	0	1
Europe	3,942 **(7,114)**	3,938 **(7,092)**	5 **(23)**
Belgium	1	0	1
Denmark	9	8	1
France	1	0	1
Germany	6	0	6
The Netherlands	1	0	1
Norway	1	0	1
Sweden	3	0	3
Switzerland	1	0	1
Uk/Ireland	163	163	0
UK	2,984	2,983	3
North America	4 **(27)**	4 **(18)**	0 **(10)**
Canada	3	0	3
USA	20	14	7
Oceania	151 **(177)**	151 **(175)**	0 **(2)**
Australia	26	24	2

10,459 (59%) of these users could not be assigned a location (either because their transactions contained no information about their location or because this location information was not specific enough). The distribution of rest of the user locations are displayed in Table. 5.2.

5.2.3 Network Analysis

The analysis conducted on the network structure was designed to mimic that of Duxbury and Haynie (2017) who used a dataset created from the public image of a DNM. This was done, in part, to help validate the analysis (by checking that the results are similar) and in part to investigate whether or not publicly available data can be used to conduct accurate network studies (by examining any discrepancies in the results). These measurements were the density, reciprocity, transitivity, and centralization of the network (Duxbury and Haynie (2017)). These measures describe key aspects about the relationships within the network and can inform a description of the interactions between buyers and vendors. The analysis was conducted using the iGraph library for Python 2.7.

First, the density of the network was calculated. This measures the number of edges in the network as a proportion of all the possible edges. A network where few of the

nodes are connected to each other has a low density and a network where each node is connected to a high proportion of the other nodes has a high density.

This prebuilt measurement calculates the number of possible edges under the assumption that any user may be connected to i.e. make purchases from, any other user. However, not all users made purchases or sold items and it is reasonable to assume that some accounts belonged to users who just wanted to buy or just wanted to sell, therefore an additional, adapted measure of density was calculated as

$$D = \frac{\#\text{edges}}{\#\text{buyers} \times \#\text{vendors}} \tag{5.2}$$

The second measure taken was the reciprocity of the graph. This measure is the proportion of edges from users i to j that are reciprocated. When the reciprocity is low, this implies that few buyers sell back to vendors they have bought from and vice versa.

The third measure taken was the transitivity of the network. This counts the number of triangles within the network as a ratio of the number of connected triplets (the number of subgraphs comprised of 3 nodes connected by 2 edges) and has the following formula:

$$\mathbf{T} = \frac{3 \cdot \text{Number of Triangles}}{\text{Number of Connected Triples}} \tag{5.3}$$

A high transitivity means that, when two users are both connected to a third they are likely to be connected also, i.e. if two buyers make purchases from the same vendor one will also make a purchase from the other.

The final measure was the centralisation of the network. This is a measure of how much the network is focused around one, or a small proportion of, specific node(s). It is calculated by taking the ratio of the difference between the maximum degree in the network and every other degree and the sum of these differences in a star network with the same number of nodes (Scott (2017); Duxbury and Haynie (2017)).

5.2.4 ERGM Analysis

Duxbury and Haynie (2017) use ERGM to identify key characteristics of vendors that likely determine the number of edges they have to buyers. They evaluated the vendors' trustworthiness (measured as the sum of ratings across all their transactions), affordability (measured as the average price of their products sold in the previous 6 months), and product diversity as predictor variables, controlling for the vendors' country of origin and buyers' degree.

Their methodology was replicated on each weekly snapshot of the network, with edges weighted by transaction value, using the iGraph package for R as the network itself was too large for this analysis.

5.3 Reddit Study

The forum data from the Reddit threads described in Section. 4.3 was used to understand how Dark Net Market users responded to the law enforcement interventions Operation Hyperion and Operation Bayonet. A combination of quantitative and qualitative research techniques were applied to the data to understand how the community interacted and what it was discussing. This subsection outlines the tools and methods used to carry out this research.

5.3.1 Measuring the Community

A network was created for each subreddit such that contributors were represented as nodes and directed edges were drawn between nodes if one had commented on the post of the other. The same measurements of reciprocity, transitivity, and centralization and the adapted measurement of density outlined in the methodology section for the study of *Silk Road 2.0* were applied to these networks using the iGraph library for Python 2.7.

Not all comments on a post are likely to be direct replies to the post – some may be replies to comments. This information – of when a comment is directly replying to a comment was not recorded by the data source and so could not be directly measured. Instead, for each subreddit, a second network was constructed such that edges were drawn between two nodes if one had commented on the other's post or if they had both commented on the same post. The network measurements were recollected to provide a range within which the true value is expected to lie.

5.3.2 Measuring the Impact of Events

The community response to law enforcement events was assessed both quantitatively, by looking at the number of comments and posts per event, and qualitatively, by applying Grounded Theory to the content of relevant posts.

5.3.2.1 Quantitative Approach

To determine if the increase or decrease in conversation after Operations Hyperion and Bayonet were disproportionate to other population fluctuations, a moving average was calculated for the number of posts, comments and contributors of each subreddit. Increases or decreases above or below 3 standard deviations from the mean were considered to be anomalous values of population change.

5.3.2.2 Qualitative Approach

In order to qualitatively assess the impact of different law enforcement interventions, relevant posts and comments were extracted from the dataset by searching for key words in posts and comments made during each event. The list of keywords and date parameters imposed on the search were initially set to try and collect all relevant posts/comments even if some irrelevant posts/comments were also collected. This returned a dataset that was then inspected to remove irrelevant posts and comments, to identify more key words that seemed associated with the incident but had not been searched for, and to determine if the observation dates needed to be expanded.

When as many relevant posts and comments as possible had been found, Grounded Theory was applied to the dataset. This approach involves an iterative process of reading the forums to identify key themes and/or concepts that are then substantiated through rereading and further data collection ensuring that any conclusions are grounded within the data (Grounded Theories Ltd (2016)).

To do this, each dataset was read through once noting reoccurring themes, topics, and hypotheses. This informed a set of categories which were then applied to the dataset using Directed Content Analysis, i.e. the posts/comments were sorted into the identified themes to produce evidence for the hypotheses.

As the analysis for each event is data driven and therefore context specific, the exact approach taken for each event is given in more detail in Chapter 8.

Chapter 6

Results of Cross Market Study

This section presents the results of the cross market study described in previous chapters. It presents the general statistics found which describe the vendors active on the ecosystem and the impact of Operation Onymous on the vendor population.

The results from this section were used to discuss several hypotheses of how users respond to different event impacts, namely hypotheses 5, 6, 7, 11, 13 and 16. These predict how users will be impacted when the ecosystem exhibits greater instability or when they lose their account, lose money, lose reputation and/or are made more aware of the presence of law enforcement. The results of this study are also used to evaluate hypothesis 1 which describes how reducing the cashflow of the ecosystem affects its overall trade.

6.1 General Statistics

6.1.1 Vendors

Of the 10,092 vendors in the database, 7,530 (75%) had accounts on only one DNM and vendors maintained accounts on an average of 1.51 different DNMs. The full distribution is displayed in Figure 6.1. This result is lower than that reported by Broséus et al. (2016) where 80% of vendors were active on only one DNM. This could be a result of that study featuring much fewer (6) DNMs.

Vendors had an average account lifetime of 91 days and a average total lifetime of 126 days. The longest running account lasted for 552 days and the vendor with the longest lifetime was active for 629 days. The full distribution of vendor lifetimes is presented Figure 6.2. This differs to the analysis found by Soska and Christin (2015) where vendors were found to be active for an average of 220 days. This could be because our study

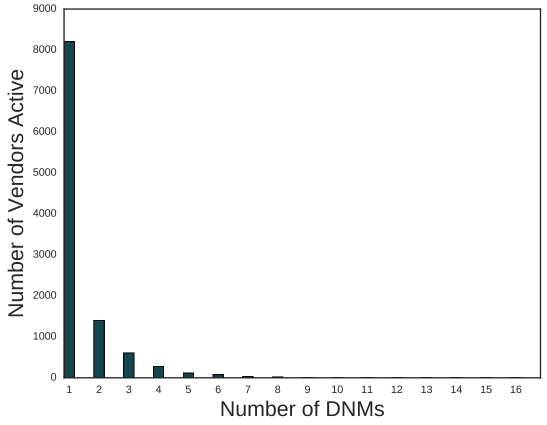

FIGURE 6.1: **Number of Vendor Accounts** This plot shows the number of vendors controlling each number of vendor accounts.

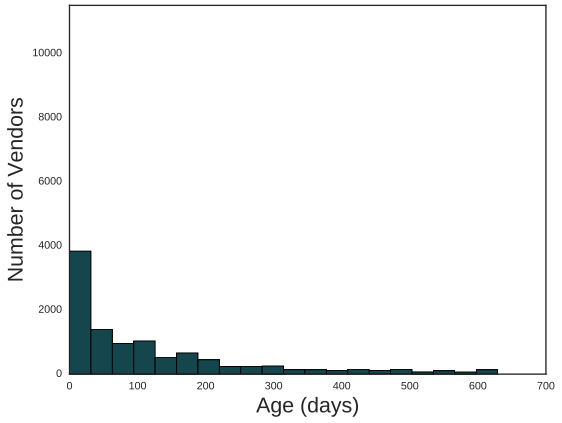

FIGURE 6.2: **Histogram of the Account Ages** This plot shows the ages of vendor accounts measured in days.

contains many smaller DNMs that were active for a shorter period of time, reducing the average.

The vendor population, measured by number of accounts, was plotted over the measurement period and displayed in Figure 6.3. This shows the total weekly population broken down into the existing number of accounts (green), the number of new accounts created by vendors who were already operating on the ecosystem (orange) and the number of new accounts created by new vendors (purple). This figure shows two large falls in population once after Operation Onymous and once after the closure of *Evolution* in an exit scam.

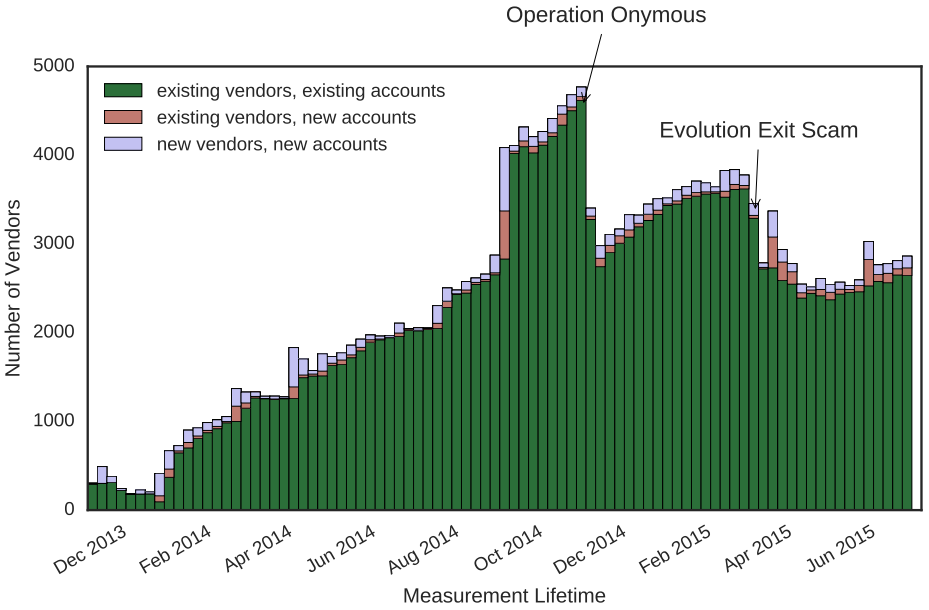

FIGURE 6.3: **Ecosystem Population** Population change over time. It shows the number of new accounts from new vendors (purple), the number of new accounts from existing vendors (orange) and the number of existing accounts (green).

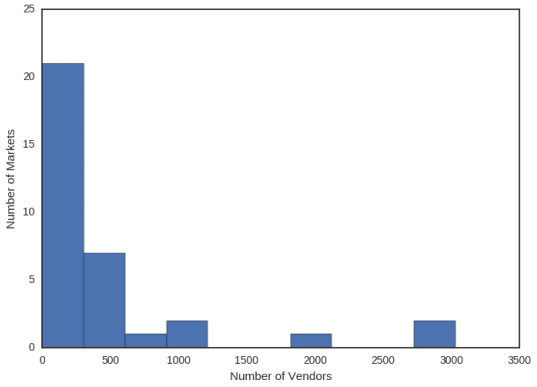

FIGURE 6.4: **CDF of DNM Sizes**

6.1.2 DNMs

The DNMs in this study had an average total vendor population (the number of vendors active during the entire measurement period) of 481 vendors. The maximum was 3,032 (*Agora*) and a minimum of 4 (*Bloomsfield*). Figure 6.4 shows a CDF plot of the DNM populations. 88% of the DNMs had a less than 500 different vendors trading during their lifetime.

The longest running DNM lasted for 552 days and the shortest lived ran for 6 days with the average being 143 days. During the measurement period, 25 DNMs closed

down. 5 were shut by law enforcement (in Operation Onymous); 18 were closed by their administrator, of these 11 closed in a suspected exit scam and the administrator stolen the coin held on the DNM at the time of closure, 1 shut in response to a hack, and 6 were closed voluntarily for other reasons, for example because they were not profitable; 2 DNMs shut down for unknown reasons.

6.2 Analysis of the Impact of Events

6.2.1 Operation Onymous

Operation Onymous had the largest impact on the vendor population during the measurement period. The population reduced from 4,596 accounts controlled by 3,593 vendors on 28 October to 2,974 accounts controlled by 2,465 vendors on 11 November 2014. After this decrease, the population begins to rise but does not reach its previous height during the measurement period. 2,302 vendors (64%) continue to trade on 4,161 accounts 1,399 of which are created after the intervention.

The vendor population of the DNMs shut down in Operation Onymous was 1,449 controlling 1,532 accounts on these DNMs and 627 on other DNMs. 1,025 of these vendors were only operating on the DNMs that were shut. Of the total vendor population, 2,391 (67%) continued to trade: 2,302 maintained 2,762 existing accounts and 834 created 1,399 new accounts. This means that 90% of the accounts that these vendors were able to maintain were maintained. Of the 2,391 vendors who continued to trade, 402 were vendors who were operating on at least one DNM that was shut down and at least one DNM that was not. In contrast, of the 1,025 vendors only operating on the affected DNMs only 75 (7%) continued to trade by creating 115 new accounts.

Further, of the vendors who were only operating on unaffected DNMs, the vast majority (89%) continue to trade after Operation Onymous. That is 1,914 of the 2,144 vendors either maintained 2,175 of their existing accounts and/or created 860 new ones. As such, the impact of Operation Onymous has been greatest on the vendor population exclusively trading on the DNMs it shut down, then the population who traded on affected and unaffected DNMs, and is least felt by the population who was not trading on the DNMs it shut.

To compare the distribution of the affected population (vendors operating at least one account on a DNM shut down in the operation) into categories Continues Trading and Does Not Continue Trading and the distribution of the unaffected population, a Chi-Square test is used. This tests the hypothesis that the unaffected population follows the

TABLE 6.1: Chi Square Statistics

	Observed	Expected
Continues Trading	1,914	477
Does Not Continue Trading	230	972

TABLE 6.2: Average Values for Vendor Characteristics Before and After Operation Onymous.

Variable	All Vendors Before Onymous	Unaffected Vendors who Continued to Trade	Affected Vendors who Continued to Trade	Vendors who Ceased Trading
Number of Accounts	1.58	1.26	2.87	1.10
Age	154	120	232	148
Rank	252	243	185	319
Number of PGP Keys	2.01	1.67	3.60	1.31
Number of Usernames	1.20	1.07	1.62	1.09

same distribution as the unaffected population. The observed results are the proportions of the unaffected population that do and do not continue to trade, and the expected results are the proportions of the affected population that do and do not continue to trade, as described in Table. 6.1. This produces a Chi-Square statistic of 4,895 with 1 degree of freedom, and a p-value of < 0.0001. It therefore cannot be concluded that the affected population and unaffected population are distributed into the categories Continues Trading and Does Not Continue Trading in a similar manner.

An examination of the population characteristics reveals further differences between the affected and unaffected populations. The vendor population active before Operation Onymous had accounts on an average of 1.58 DNMs and had been operating for an average of 154 days. Further, they had an average ranking of 252, calculated using the method described in Section 5.1.4. Similarly, the unaffected vendors who continued to trade had accounts on, on average, 1.26 different DNMs, an average lifetime of 120 days, and an average ranking of 243. In contrast, the affected vendors who continued to trade had accounts on an average of 2.87 DNMs, an average lifetime of 232 days and an average ranking of 185, i.e. had more accounts, were older and more highly ranked. A summary of these values, and additional characteristics, is given in Table. 6.2.

To evaluate if any of these characteristics were significant predictors of the type of vendor who continued to trade after Operation Onymous, logistic regression analysis was applied to the results using the sci-kit learn module for Python 2.7[1].

[1] http://scikit-learn.org/stable/

TABLE 6.3: The Variance Inflation Factor of each variable in the logistic regression model.

Variable	VIF
Age of Vendor	1.8
Number of Accounts	4.2
Average Rank	21.3
Max Rank	18.4
Min Rank	53.5
Number of PGP Keys	2.5
Number of Usernames	1.6
Number of Markets Closed	2.7
Affected by OO	1.3

Of the 3,928 vendors for whom each variable could be collected, 2,515 continued to trade and 1,413 did not. As can be seen in Figure 6.4, each of the variables was significant except for the number of PGP keys employed by the vendor. The model had an accuracy of 96% (the mean result of 10-fold cross validation was also 96%).

To measure the collinearity of the variables in the model, the prebuilt function variance_inflation_factor from the Python package statsmodels was used. This returned a high VIF score for the three measures of rank (as would be expected) and sufficiently low scores for the remaining variables. The VIF scores for each variable is given in table 6.3.

The model showed that the age of the vendor negatively influenced whether or not they would continue to trade, i.e. the older they were the less likely they would. The number of accounts a vendor had also had a negative influence on the probability the vendor would continue to trade, though to a much greater degree. Interestingly, the number of usernames the vendor used positively increased the chances of the vendor continuing to trade, this perhaps implies that vendors are more likely to continue trading if they have accounts they are working to make unlinkable operating on different markets, which seems intuitive if vendors are concerned about the information law enforcement has collected on them during Operation Onymous.

The average rank value has a negative influence on the likelihood that a vendor will continue to trade, this means that vendors with higher reputations are more likely to continue to trade. This effect is also seen with the maximum rank value that vendors had earned during their lifetimes. The minimum rank value that they had earned has a positive influence, i.e. the lower the minimum reputation a vendor received, the less likely they are to continue trading.

The number of markets a vendor has seen closed during their lifetime positively influences if they choose to continue to trade. This means that vendors who have seen more DNMs

TABLE 6.4: The Influence of the Variables in the Model Built to Predict if a Vendor Will Leave After Operation Onymous.

	Coefficient	Standard Error	z	$P > \|z\|$
Age of Vendor	-0.0242	0.002	-15.5	0.000
Number of Accounts	-2.93	0.240	-12.2	0.000
Average Rank	-0.0095	0.001	-7.29	0.000
Max Rank	-0.0115	0.002	-6.35	0.000
Min Rank	0.0210	0.003	6.98	0.000
Number of PGP Keys	-0.159	0.175	-0.909	0.363
Number of Usernames	1.390	0.251	5.54	0.000
Number of Markets Closed	0.711	0.033	21.8	0.00
Affected by OO	-2.59	0.237	-11.0	0.000

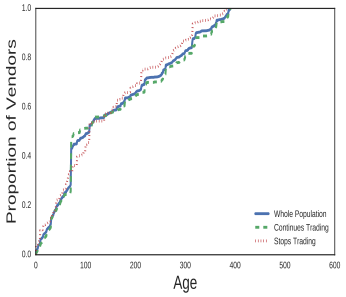

FIGURE 6.5: **CDF plot of the ages of the vendors** who were active during Operation Onymous.

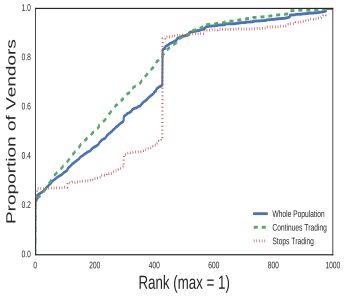

FIGURE 6.6: **CDF plot of the ranks of the vendors** who were active during Operation Onymous.

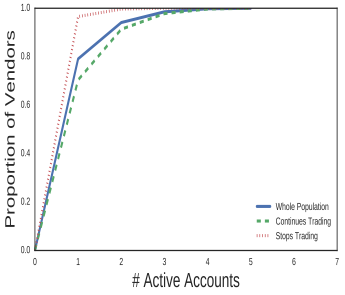

FIGURE 6.7: **CDF plot of the number of accounts** owned by the vendors who were active during Operation Onymous.

close are more likely to continue to trade, potentially because they are more accustomed to the disruption of the ecosystem.

To more concretely understand how each variable may have affected the vendor population, the CDF of the age, rank and number of accounts owned by the vendors active during Operation Onymous was plotted next to the CDFs for the vendors who continued to trade and those who stopped trading after the operation. Figure 6.5 shows that the age profile of vendors who stopped trading compared to those who continued was fairly similar. However, Figure 6.6 shows that a much lower proportion of vendors who stopped trading had lower ranks (higher reputations) which explains why their reputations were, on average, lower. Finally, Figure 6.7 shows that a much greater proportion of the vendors who continued to trade only had one account before the operation.

6.2.2 Evolution Exit Scam

The results of Operation Onymous were compared with the impact that the exit scam of the DNM *Evolution*, which took place on 17 March 2015, had on the population. This event was chosen because it was the only other event to have a comparable impact on the population (see Figure 6.3)[2].

During this event, the vendor population fell from 3,911 accounts controlled by 3,197 vendors to 2,916 accounts controlled by 2,505 vendors. The fall in the population is observed to take place slightly before to the assumed closure of the site, this is potentially because, prior to closing, the admin team of *Evolution* blocked users from withdrawing their funds, an action which may have warned users to leave the site.

The reduction in vendor and account population (25% and 22% respectively) is much smaller than the one that occurred after Operation Onymous. Once again, the Chi-Square test was used to test the significance. It produced a Chi–Square test statistic of 166 at a p–value of $5.06 \cdot e^{-38}$ and so the null hypothesis that the two event populations exhibited a decrease of the same proportion should be rejected.

After Operation Onymous, the observed population begins to increase (see Figure 6.3). It does not reach pre-Onymous figures, perhaps because the Evolution Exit Scam occurs. After *Evolution* closes the population is boosted by an influx of new accounts and new users, however it then continues to decrease further before starting to increase.

It would be expected that, as Operation Onymous appears to have a bigger impact on the ecosystem population, users collectively lost more financially in Operation Onymous when compared to the Evolution Exit Scam. This would be because a financial loss in a DNM closure impedes a user's ability to continue trading.

However, the Evolution Exit Scam is rumoured to be worth more than $12 million (Krebs (2015); DeepDotWeb (2015d)) or even as much as $34 million (Chung (2015)) whereas Europol announced that just $1 million was seized in Operation Onymous (Europol (2014)). If these figures are accurate then the ecosystem lost more in the Evolution Exit Scam than Operation Onymous. Calculating the average loss per vendor (using the population sizes in the dataset and assuming only vendors lost money) gives approximately $690 per vendor after Operation Onymous and between $7,190 and $20,400 per vendor active on *Evolution* when it closed. Even with the margin of error, the loss from the Evolution Exit Scam was an order of magnitude greater.

[2]It was the only comparable event during the measurement period. Other DNMs, e.g. Sheep, have closed in similarly dramatic fashions but are out of scope for this research because they are not included in the dataset.

The administrators of *Evolution* were better placed to seize more money than law enforcement as they had the ability to prevent users from withdrawing funds in the run up to their exit. Despite this, the impact of the Evolution Exit Scam was measurably smaller than Operation Onymous.

To more directly compare the two events, the population was broken down into those affected and unaffected and the rates of those who continued to trade were compared.

2,232 (70%) of the overall population continued to trade on 3,355 (86%) accounts after *Evolution* closed. This is a similar percentage of vendors but smaller percentage of accounts than after Operation Onymous, potentially due to there being less non affected DNMs after this event. However, marginally more vendors (25%) began trading on 1,201 new accounts.

1,670 (52%) vendors had an account on *Evolution*. Of these, 491 (29%) maintained 619 of the 1,100 (56%) accounts they had on other DNMs and 521 (31%) created 797 new ones. 1,181 (71%) vendors operated only on *Evolution* and 305 (26%) of those created 464 new accounts on other DNMs, a much higher percentage than after Operation Onymous.

By contrast, the percentage (94%) of non-*Evolution* accounts owned by vendors who weren't trading on *Evolution* that were maintained was similar to the number of accounts maintained by vendors that were not affected by Operation Onymous. In addition, 269 (18%) made new accounts which is also similar to the proportion of the unaffected population who behaved in this way after Operation Onymous. This implies that the impact of the Evolution Exit Scam was also contained to those who had accounts on the affected DNM.

To see if the same characteristics were held by vendors who continued to trade after the Evolution Exit Scam as after Operation Onymous, a logistic regression model was built. There were 3,373 vendors for whom the same variables in Section 6.2.1 could be measured. 2,406 of these vendors continued to trade and 967 did not. The results are given in Table. 6.6. They show that the number of usernames they used did not significantly predict if a vendor would continue to trade, whereas the other variables did. The model had an accuracy of 91% (the mean result of 10-fold cross validation was also 91%).

The collinearity of the variables was again measured. The VIF scores for each variable is given in table 6.5. Once again, the VIF values for the rank variables are too high. However, for this model, so is the VIF value for the number of accounts[3].

[3]These values are included in the thesis at the request of a viva examiner and so the variables with high VIF scores have not been removed from the model.

TABLE 6.5: The Variance Inflation Factor of each variable in the logistic regression model.

Variable	VIF
Age of Vendor	2.6
Number of Accounts	6.5
Average Rank	15.5
Max Rank	13.3
Min Rank	35.5
Number of PGP Keys	3.0
Number of Usernames	3.3
Affected by Exit Scam	1.3
Active During OO	3.0

TABLE 6.6: The Influence of the Variables in the Model Built to Predict if a Vendor Will Leave After the Evolution Exit Scam.

	Coefficient	Standard Error	z	$P > \|z\|$
Age of Vendor	-0.0122	0.001	-13.7	0.000
Number of Accounts	-0.896	0.181	-4.95	0.000
Average Rank	-0.0073	0.001	-7.46	0.000
Max Rank	-0.0088	0.001	-6.46	0.000
Min Rank	0.0170	0.002	7.36	0.000
Number of PGP Keys	2.14	0.264	8.11	0.000
Number of Usernames	-0.033	0.301	-0.110	0.913
Affected by Exit Scam	-3.14	0.145	-21.6	0.000
Active During OO	-1.27	0.277	-4.59	0.000

The logistic regression model describes the vendors who continued to trade as having, on average, fewer accounts, a lower maximum, and average ranking (and so higher minimum, and average reputation), more PGP keys and experiencing over 3 times as many market closures.

An additional variable that was considered in this model was whether or not the vendors affected by the Evolution Exit Scam were also trading during Operation Onymous. It was found that vendors who stopped trading were more likely to have lost an account in the Evolution Exit Scam and more likely to have been operating during Operation Onymous.

The CDF of vendor ages is shown in Figure 6.8. Here it can be seen that a greater proportion of vendors who stopped trading had accounts that were less than 250 days old than those who continued trading and the general population active at the time. The fact vendor age was not a significant predictor, however, could be because operating for longer on the ecosystem strengthens the resilience of some vendors but weakens it in others.

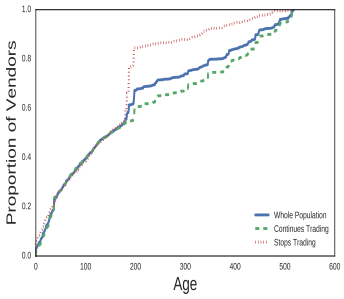

FIGURE 6.8: **CDF plot of the ages of the vendors** who were active during the Evolution Exit Scam.

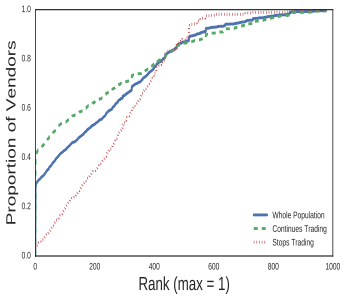

FIGURE 6.9: **CDF plot of the ranks of the vendors** who were active during the Evolution Exit Scam.

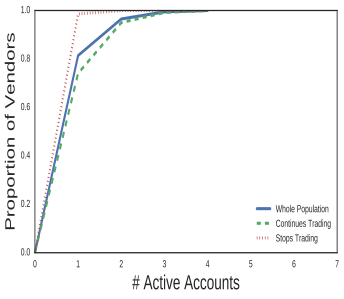

FIGURE 6.10: **CDF plot of the number of accounts** owned by the vendors who were active during the Evolution Exit Scam.

An examination of the CDF plot of the rank of vendors (Fig. 6.9) shows that a smaller proportion of vendors who stopped trading had smaller ranks (higher reputations) and a greater proportion had higher ranks (lower reputations). This explains why vendors who continued to trade had a lower average ranking. Figure 6.10 is the CDF plot of the number of accounts owned by vendors, this shows that a much bigger proportion of the vendors who stopped trading only had one account and the maximum number of accounts owned by the vendors who stopped trading was less than that of the population of vendors who continued to trade.

A direct comparison between the vendors who continued to trade after Operation Onymous and after the Evolution Exit Scam shows they shared some, but not all, characteristics. The vendors who continued to trade after the Evolution Exit Scam had a similar number of accounts and PGP keys as those who continued to trade after Operation Onymous. However, they had lower rank scores (higher reputations) and were more likely to have been affected by the event. The results are summarised in Table. 6.7.

As the vendors operating during the Evolution Exit Scam can be older than those operation during Operation Onymous, they may have accumulated a higher reputation and this could explain why the vendors who continued to trade after the Evolution Exit Scam had, on average, lower ranks. However, another explanation is that, if the ecosystem had become more competitive, vendors may have only been able to transfer their business to a new DNM if they had a higher reputation.

As for vendors who continued to trade after the Evolution Exit Scam being more likely to have been affected, this could indicate that, as more large events hit the ecosystem, vendors became less worried about large DNM closures. Alternatively, the fact that the DNMs in Operation Onymous were closed by law enforcement, potentially resulting in

TABLE 6.7: Comparison of Vendors who Continued to Trade After the Evolution Exit Scam and Operation Onymous.

Variable	Operation Onymous	Evolution Exit Scam
Number of Accounts	1.76	1.89
Average Rank	371	334
Max Rank	184	134
Min Rank	260	214
Number of PGP Keys	1.71	1.73
Affected by Event (%)	35	46

TABLE 6.8: DNMs from the Observation Period that Closed Amicably

DNM	Dates in Operation	Lifetime (days)	Total Population	Closure in Observation Period
1776	19/4/2014 - 2/10/2014	166	37	No
Agora	3/12/2014 - 6/9/2015	277	3,032	No
Darkaby	30/1/2014 - 1/5/2014	91	789	Yes (in dataset until 12/5/2018)
Freebay	2/12/2013 - 28/2/2014	88	174	No
Panacea	27/10/2014 - 13/2/2015	109	21	Yes

sensitive information about vendors and buyers being seized, could have been a greater deterrent to vendors to continue trading.

6.2.3 DNM Closure

DNMs can close amicably as well as being shut by law enforcement or closing in an exit scam. In amicable closures, users are often warned and given time to remove any cryptocurrency held in escrow or put in place other measures to minimise the harm of the closure. It therefore seems reasonable that such a closure would have a much smaller impact on the ecosystem than either Operation Onymous or the Evolution Exit Scam. To measure this, a comparable amicable closure was found and its impact measured.

7 of the DNMs in the dataset closed amicably (rather than being shut down, being attacked before closing, or leaving in an exit scam). Of those, 3 closed during the measurement period: *Darkbay, Panacea*, and *Underground Market*. *Evolution* had a total population of 3,000 vendors (within the dataset) which is much larger than any of these markets that had populations of 789, 21, and 22 respectively. *Darkbay* was used to compare an exit scam to an amicable closure as it was the closest in size to *Evolution*. Other DNMs, of more comparable size, have closed amicably however Darkbay was the most suitable for which data was available. More details on each of the DNMs and their suitability is given in Table. 6.8.

The population prior to *Darkbay's* closure (taken between 23 April and 13 May 2015) was 1,674 in control of 1,993 accounts. Of these, 157 (9%) users had accounts on *Darkbay* and 151 were only trading on *Darkbay*. 10 (6%) of the vendors who had an account on *Darkbay* continued to trade 9 created new accounts and 4 maintained existing accounts. 4 (3%) of the vendors who only had an account on *Darkbay* continued to trade and did so by creating 4 new accounts. By contrast, 93% of the population not trading on *Darkbay* continued to trade.

When comparing to the impact of the Evolution Exit Scam, the closure of *Darkbay* seems to have had an even bigger impact on its affected population than *Evolution* does with the impact on the unaffected population in each event approximately the same. This could be because vendors on *Evolution* were unable to retrieve any coin they had in escrow on the site necessitating them to continue trading in order to recuperate their costs. It could also be because vendors on *Darkbay* were less able to gain access to other DNMs as vendors on *Evolution*.

When comparing to Operation Onymous, the impact of *Darkbay's* closure on the overall ecosystem was much smaller, presumably because the marketplace was much smaller than the combined size of all the marketplaces closed in Operation Onymous. Additionally, the impact on the unaffected population was about the same. However, a larger proportion of vendors who had at least one account on a DNM closed in Operation Onymous and at least one account on a DNM that was not closed continued to trade after Operation Onymous than in the equivalent population when *Darkbay* shut. Similarly, a larger proportion of the vendors only trading on DNMs shut during Operation Onymous continued to trade than the proportion of *Darkbay* only vendors who stopped trading when it closed.

6.2.4 Vendor Movement

As well as looking at who stopped and continued trading, it was of interest to examine where the vendors who continued to trade created new accounts. After Operation Onymous, vendors could have chosen their new DNM(s) at random or been attracted by certain characteristics, such as the DNM size or age.

It was first tested to see if the vendors moved to new DNMs at random. This was measured using the Chi-square test which compared the distribution of vendor accounts to the uniform distribution. The test returned a test statistic of 2,660 (p-value < 0.0001), causing the hypothesis that the DNMs were chosen at random to be rejected. Notably, the most popular DNM was *Evolution* which received 44% of the new vendor accounts.

TABLE 6.9: Results of ERGM analysis evaluating the variables of the size and age of DNMs in the ecosystem on the number of new DNM accounts made post Operation Onymous.

Variable	Estimate	Standard Error	p–value
Edges	-1.16	0.653	0.0777
Age	0.000114	0.00157	0.942
Size	0.000954	0.000300	0.00185

The influence of the age and size of the DNMs on the decision by vendors on where to move was measured using ERGM analysis. The ecosystem was modelled as a network such that the DNMs were represented by nodes and directed edges represented the movement of vendors (i.e. an edge with weight one drawn from one DNM to another represents one vendor losing an account on the first DNM and creating an account on the second). This is depicted in Figure 7.12.

The ERGM analysis was conducted using the statnet package for R[4] and the results are given in Table. 6.9.

The model shows that the age of the DNM was not a statistically significant variable (p–value = 0.94234), however the size of the DNM was. The baseline probability of a vendor creating an account on a new DNM, given the DNM they were previously active on, was 0.238. This increases to 0.239 when the size of the new DNM is increased by one vendor.

As can be seen in Figure 7.12 *Evolution* was the largest DNM at the time, though comparable in size to *Agora*. Other variables that could not be measured were the reliability of the DNM (e.g. the DNM's "uptime" and functionality), the variety of products available, where the buyers appeared to be moving to, and community consensus.

6.2.5 Summary

An analysis of the impact of Operation Onymous on vendors shows that the affects were greatest on those vendors who were solely operating on the DNMs shut in the operation. Further, the vendors most likely to cease trading were those with lower reputations and less trading accounts. It can therefore be argued that Operation Onymous did not have a universal impact on the vendor population, nor did it target what are likely to be the largest vendors.

Operation Onymous did have a measurably bigger impact than the Evolution Exit Scam. This was seen both in terms of the reduction in the size of the ecosystem and in the

[4] https://cran.r-project.org/web/packages/statnet/index.html

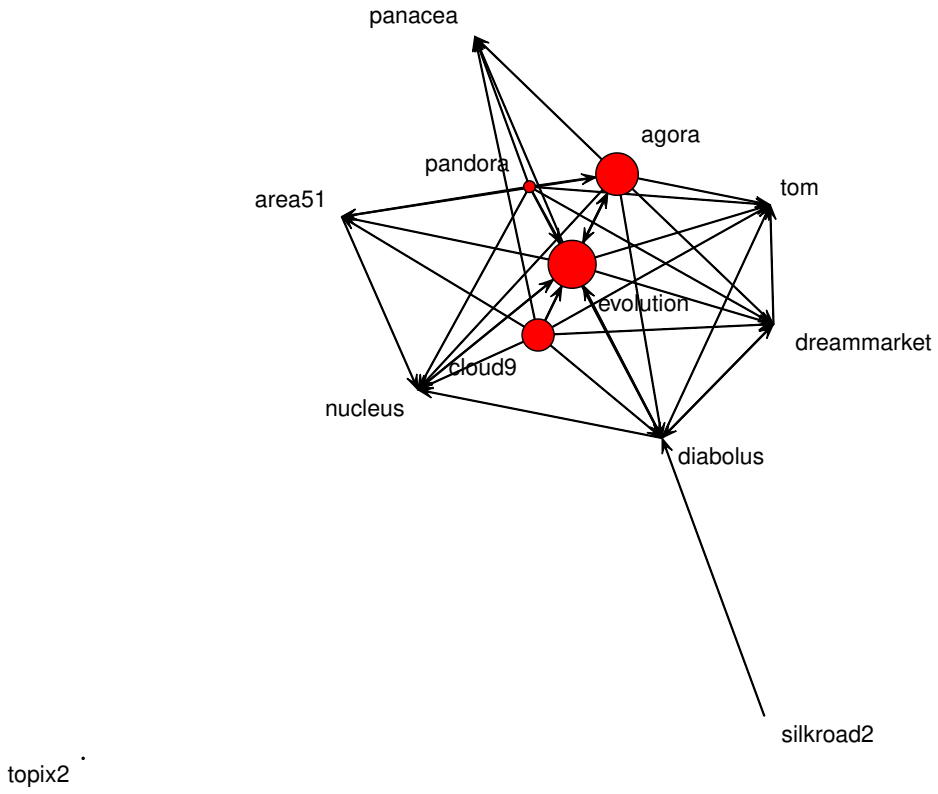

FIGURE 6.11: **Movement in Ecosystem Post Operation Onymous** This plot shows the DNMs active before and after Operation Onymous represented as nodes such that the size of the node is proportional to the number of vendors active on the DNM and the movement of vendors after Operation Onymous represented by directed edges.

fact that more affected vendors continued to trade after *Evolution* closed. This could be because vendors were more concerned by a law enforcement effort than a market closure.

Examining where vendors moved to after Operation Onymous showed that the DNMs that vendors chose to create new accounts on were likely not selected at random, however the age of the DNM did not seem to be a factor taken into consideration.

Chapter 7

Silk Road 2.0 Results

This section presents the results of the analysis on the *Silk Road 2.0* dataset. It begins with the general statistics found about the DNM then discusses how the network changes over time, and how buyers and vendors interacted on the site.

Analysis of these results has been used to answer hypotheses 12, 14 and 19. These hypotheses describe the comparative impact between losing an account and losing coin on a user's resilience, how losing reputation might affect their relationships with their trading partners and how, in turn, losing a trading partner might affect a user's resilience.

7.1 General Statistics

There are 42,169 user nodes in the network, connected by 190,802 transaction edges between 110,669 buyer/vendor pairs. The network contains 909 (2%) vendors (142 of which are labelled as UK based), 33,123 (79%) buyers (16,086 are labelled as UK based), and 8,260 (20%) isolated nodes who did not participate in any transactions. Figure 7.1 describes the network with nodes coloured for user type and figure 7.2 describes the network with nodes coloured depending on if they are labelled as UK based or not.

The average indegrees and outdegrees of nodes, as well as general buyer and vendor characteristics are described in Table. 7.1. Broadly, buyers were involved in many fewer purchases than vendors (by two orders of magnitude), and kept accounts for, on average, half the amount of time as vendors. This is consistent with a description of buyers being more superficially engaged with the site than vendors.

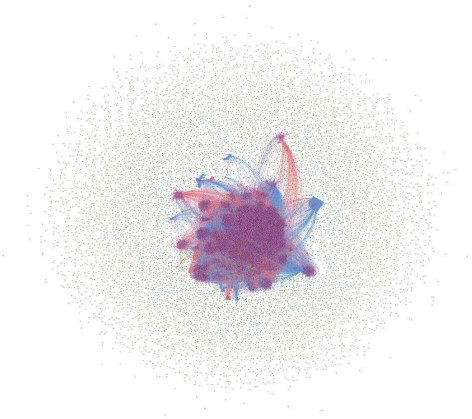

FIGURE 7.1: **Image of Network Coded by Type of Node.** This figure shows the network with buyer nodes coloured purple, vendor nodes coloured red, users that were both vendors and buyers coloured blue, and isolate nodes coloured green.

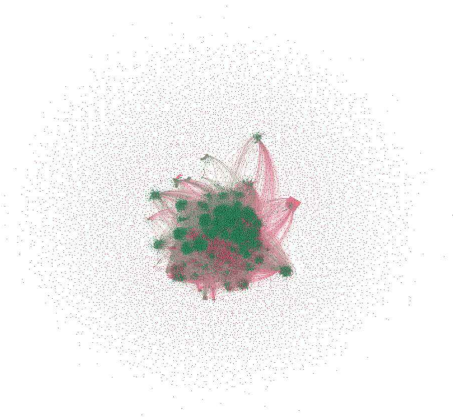

FIGURE 7.2: **Image of Network.** This figure shows the network with nodes that have full profiles in the dataset coloured red and those that do not coloured green.

7.1.1 Network Analysis

First, the measurements taken by Duxbury and Haynie (2017), namely the density, reciprocity, transitivity, and centralisation of the network are repeated. These measures describe key aspects about the relationships within the network. The density of a network reveals how many of the possible connections have been made. If the network has a high density, this implies that buyers are each making a several purchases whereas a low density implies that buyers only tended to make a few purchases from a limited set of vendors. Reciprocity and transitivity can be used to measure if it is common behaviour for vendors to also make purchases and buyers to also sell products and, crucially, if a buyer purchasing from a vendor makes them more likely to also sell to that vendor. Finally, the centralisation of the network measures the extent to which activity congregates around a limited set of vendors, i.e. if there are vendors who conduct the majority of sales or if sales are more equally distributed across the vendor population.

As these measures have been used by Duxbury and Haynie (2017), the repetition of their analysis allows for a comparison of the two networks studied. Duxbury and Haynie (2017) apply these network measures to *Cryptomarket*, a much smaller network than *Silk Road 2.0*. Differences in the results of this analysis may imply that smaller DNMs form different types of networks to larger DNMs, for example they might be more dense because there are fewer users who do not trade regularly or they may be more likely to be centred around popular individuals. Duxbury and Haynie (2017) reconstructed

the *Cryptomarket* network from a publicly available dataset and comparing their results to those from this study may also help to determine if their methodology creates an accurate image of the network or affects these measures.

The network measures are presented in table 7.1.

The network's density was calculated to be 0.000107 (to 3 significant figures) reflecting a low interconnectedness of the network. This value shows that only a very small fraction, 0.01% of the possible vendor and buyer combinations that could have been made, were made. This may in part be explained by the fact that many of the buyers did not make any purchases at all but also implies that buyers tended to buy from a small subset of vendors, instead of a large variety.

Whilst the density found by Duxbury and Haynie (2017) was also small, it was an order of magnitude greater than the density of this network. This is despite the fact that both the outdegree of buyers and indegree of vendors were much lower. However, this network contained a much higher proportion of isolates which would decrease the density. In addition, a higher proportion of the users in the network of Duxbury and Haynie (2017) were vendors which meant that a greater proportion of the edges could be formed between nodes. Indeed, if the density of the network was measured as the proportion of edges that were formed out of the set of potential edges from buyers to vendors it would be the much higher value of 0.00634.

The second measure taken was the reciprocity of the graph. The reciprocity of the network was very low, 0.0000629, as one might expect, and indicates that vendors did not tend to make purchases from the buyers that bought from them. Interestingly, however, this did happen for one pairing. One UK user, labelled as a buyer, who made 11 purchases (of Meth and an unknown product) and 2,258 sales (selling a variety of different products), sold to and bought from a non-UK based user 6 and 11 times, respectively. Duxbury and Haynie (2017) found a reciprocity of 0.

The third measure taken was the transitivity of the network. The transitivity of the network was $5.63 \cdot \exp^{-5}$, another measure expected to be low as a high transitivity relies on buyers also being vendors in the network. Duxbury and Haynie (2017) found the transitivity to be 0. The properties of reciprocity and transitivity are not particularly useful measures for these networks as, in most instances (>99%), accounts were used for buying or vending and not both, limiting the amount of reciprocal activity.

Finally, I looked at the centralisation of the network. Both the outdegree and indegree centralisations of the network were low (0.0107 and 0.230 respectively). This could be a factor of the number of isolates in the graph, indeed removing these increases both the outdegree and indegree centralisation (to 0.0133 and 0.360 respectively). Duxbury and

Haynie (2017) also found that the outdegree centralisation was lower than the indegree centralisation which implies that sales are more likely to be concentrated onto specific vendors than buyers.

Whilst the measures taken support similar interpretations of the network, there are some differences in the measurement values, with some measurements differing in orders of magnitude. These disparities could be a result of the data used. It could be the case that *Cryptomarket*, the DNM analysed by Duxbury and Haynie (2017) exhibited slightly different behaviour to *Silk Road 2.0*, potentially as a result of an overall evolution of ecosystem behaviour as the DNMs were not operational at the same time.

Alternatively, it could be due to the fact that the data used by Duxbury and Haynie (2017) is not restricted by country, whereas the data used in this study contains complete information on UK based users only. As the dataset does not contain the transactions between non-UK based users their network degrees will be smaller. This would have the potential to increase the centralisation value and decrease the density. Potentially, a similar effect may have occurred in the analysis of Duxbury and Haynie (2017) as their data was also restricted (to product type). The lower centralisation value, however, perhaps indicates that the sale of opioids in *Cryptomarket* was more of a self-contained community than UK based users in *Silk Road 2.0* (because there are less edges missing).

Another possible explanation for the differences is the impact of the different ways the networks were constructed. Duxbury and Haynie (2017) create a network which represents the publicly available reviews on the DNM *Cryptomarket*. Whilst reviews for *Cryptomarket* were compulsory, if it was possible to leave one review when multiple items were purchased simultaneously from the same vendor, then the number of transactions observed would be smaller than the actual number that took place. In contrast, as the data in this study is taken from the *Silk Road 2.0* server, each item will have been recorded individually, even when purchased together.

Indeed 6,694 (3.51%) of the purchases were found to have the same time stamp and review as at least one other purchase. These duplicates were removed and the tests were rerun. In this new network, the density was 0.000104, i.e., smaller, but relatively unchanged than for the previous network. The outdegree centralisation reduced to 0.00585 (from 0.0107) and so became the same order of magnitude as found by Duxbury and Haynie (2017). The indegree centralisation also reduced slightly, though only by a fraction (from 0.290 to 0.274) becoming marginally closer to the value calculated by Duxbury and Haynie (2017) (0.201).

To see how the chosen network construction affected the measurements taken, a second network was constructed which only included buyers and vendors based in the UK and

the edges between them (i.e. a network where the edges were complete, not the nodes). This network consisted of 24,393 nodes and 96,558 edges.

7.1.2 UK Only Network Analysis

10,264 of the nodes were isolates [1], i.e. some users only traded with users outside the UK. The node with the highest degree had a degree of 9,961, lower than in the full network as would be expected. The density of the network was 0.000162 and so this restriction of the nodes increased the density. The reciprocity was 0.0. The outdegree centralisation was 0.0181 and the indegree centralisation was 0.408, i.e. restricting the network representation to UK users and transactions increased the extent to which the network organised around a select number of important nodes.

Removing non UK users reduced the proportion of buyers who purchased more than once, from more than one vendor or from different vendors on the same day. This implies that many of the buyers purchased from in and outside of the UK, and that more than half of the buyers who made multiple purchases from more than one vendor on the same day did so from vendors in different countries. This could be because they were making purchases of different types of products which were only available in other countries or were attractive in the shops of non UK vendors.

The average age of buyers who bought from UK vendors was higher than those who also or exclusively made purchases with non UK vendors. This could be because they had better experiences (e.g. because their purchases were less likely to be seized in customs) or because they were less likely to be a short term buyer.

The feedback score for UK vendors calculated based on their UK transactions was lower than that for non UK vendors, the minimum value was also much lower than that for all vendors. This is surprising if purchases are more likely to be seized across borders and if this results in worse feedback. However, it could be a result of cultural differences, e.g. if non-UK buyers give higher ratings. Vendors also made less on average from their UK purchases, this could be related to the fact that these vendors receive lower ratings and aren't able to charge as much or, perhaps, they charge less for domestic UK transactions e.g. because of lower postage costs.

All the general statistics of this network are presented in Table. 7.1.

[1] An isolate is a node which is not connected to any other in the network.

TABLE 7.1: General Network Statistics

	Whole Network	UK Only
Global Network Statistics		
Total Actors	42,169	24,393
Total Vendors	909 (139 labelled as vendors, 3 labelled as buyers)	140
Indegree	Mean: 4.52, s.d. 124, range: 0 - 12,212	Mean: 3.96, s.d. 113, range: 0 - 9,961
Outdegree	Mean 4.52, s.d. 9.15, range: 0 - 455	Mean: 3.96, s.d. 8.54, range: 0 - 446
Isolates	8,260	10,264
Total Buyers	33,123 (105 labelled as vendors, 15,981 labelled as buyers)	14,006
Total Edges	190,802 (110,669 if weighted)	96,558 (51,598)
Density	0.000107	0.000162
Reciprocity	0.0000629	0
Transitivity	0.0000563	0.00003.52
Outdegree Centralisation	0.107	0.0181
Indegree Centralisation	0.230	0.408
Buyer Characteristics (UK based users who made purchases)		
Outdegree	Mean: 9.20, s.d. 12.9, range: 1 - 455	Mean: 6.89, s.d. 10.3, range: 1 - 446
Buyers who have purchased from more than one vendor	12,496 (78%)	9,792 (70%)
Buyers who have purchased from more than one vendor on the same day	8,903 (55%)	1,922 (14%)
Buyers who have purchased more than once	13,341 (83%)	10,801 (77%)
Average Feedback Score	Mean: 4.76, s.d. 0.579, range: 1 - 5	Mean: 4.08, s.d. 1.78, range: 0 - 5
Average Spent in Transactions ($)	Mean: 87.8, s.d. 149, range: 0 - 8,280	Mean: 76.5 s.d. 95.9, range: 0 - 4,890
Average Age (days)	Mean: 119, s.d. 113, range 0 - 366	Mean: 156, s.d. 116, range: 0 - 364
(Length Active)	(Mean: 61.4, s.d. 95.8, range: 0 - 357)	(Mean: 84.8, s.d. 102, range: 0 - 356)
Vendor Characteristics (UK based users who sold things)		
Indegree	Mean: 982, s.d. 1,800, range: 1 - 12,212	Mean: 690, s.d. 1,320, range: 1 - 9,961
Average Feedback Score	Mean: 4.73, s.d. 0.386, range 2.86 - 5	Mean: 4.50, s.d. 1.10, range: 0 - 5
Average Cost of Transactions with vendor ($)	Mean: 116, s.d. 393, range: 0 - 4,210	Mean: 97.2, s.d. 253, range: 0 - 2,190
Average Age (days)	Mean: 227, s.d. 119, range: 0 - 366	Mean: 233, s.d. 118, range: 2 - 366
(Length Active)	(Mean: 196, s.d. 124, range: 0 - 358)	(Mean: 182, s.d. 125, range: 0 - 358)

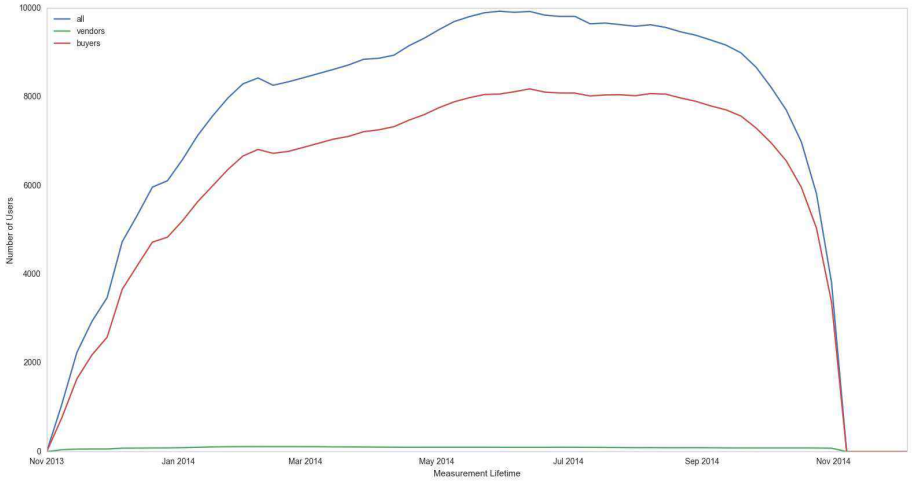

FIGURE 7.3: **Population** This plot shows the number of users, buyers, and vendors in each week.

7.2 Temporal Analysis

Analysing the way the network changes over time, both in size and in terms of its characteristics, can help to understand how the users respond to events and how different types of users interact.

7.2.1 How Does the Population Respond to Events?

Figure 7.3 shows that the buyer population grew for the first 6 months before beginning to decline, at first gradually and then rapidly in the month of its closure. The vendor population appears more stable but a closer examination reveals a sharp increase until March 2014 and then a steady decline which isn't immediately explained by the timeline (see Appendix F). There are several events that could have contributed, however, such as a hack of the site in February and the creation of multiple other DNMs in March and May which could have created competition and drawn new vendors away.

The number of transactions, and their value over time, follow similar patterns to the overall population, to begin with, however more dramatic reductions in these measures (than population decreases) occur at the end of 2013 and in February 2014. This can be seen in figure 7.4 and figure 7.5. In figure 7.6, the change in population can be directly compared to the change in the value of sales. This graph implies that, during adverse events that cause users to close their accounts, the impact of the event can be seen to a greater extent in reduction of amount spent, rather than the reduction in accounts.

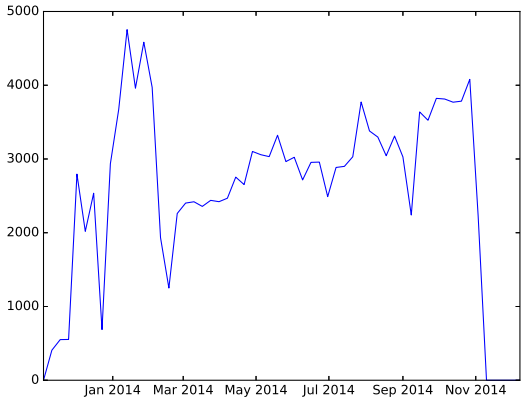

FIGURE 7.4: **Number of Transactions** This plot shows the number of transactions that take place each week.

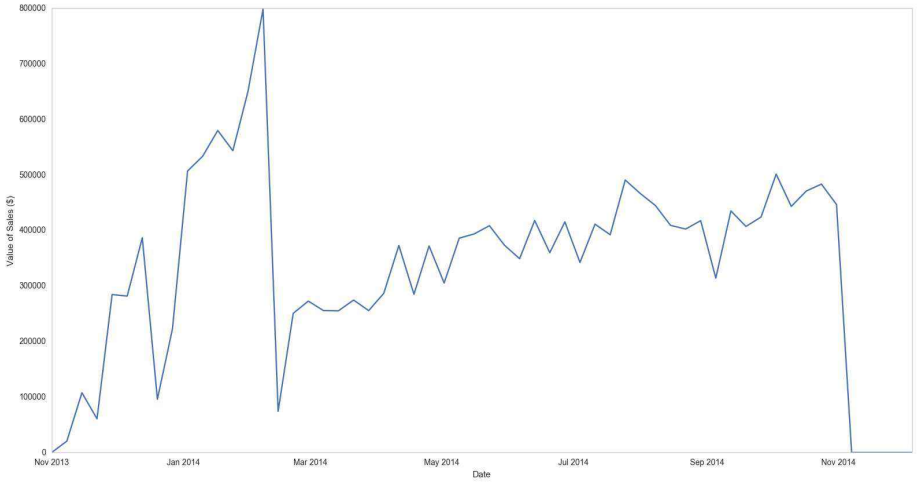

FIGURE 7.5: **Amount of Sales** This plot shows the amount, measured in USD, of sales each week.

On 9 December 2013, *Silk Road 2.0* (along with *Pandora* and *Tormarket*) suffered a DDoS attack (Digital Citizens Alliance (2013)). On 13 February 2014 *Silk Road 2.0* announced that it had been hacked and, on 15 February it shut down for 36 hours (Greenberg (2014e); DeepDotWeb (2014e)). These events could explain the reduction in the number of transactions, and subsequent population decreases.

To understand if any of the visible increases or decreases in market size could be considered significant the change in population was measured over time. This relies on the premise that, whilst the size of the DNM is expected to constantly fluctuate, if the population was not affected by events that either drove away users or deterred them from

FIGURE 7.6: **Amount of Sales and Change in Population** This plot shows the amount, measured in USD, of sales each week and the population over time plotted against each other for comparison.

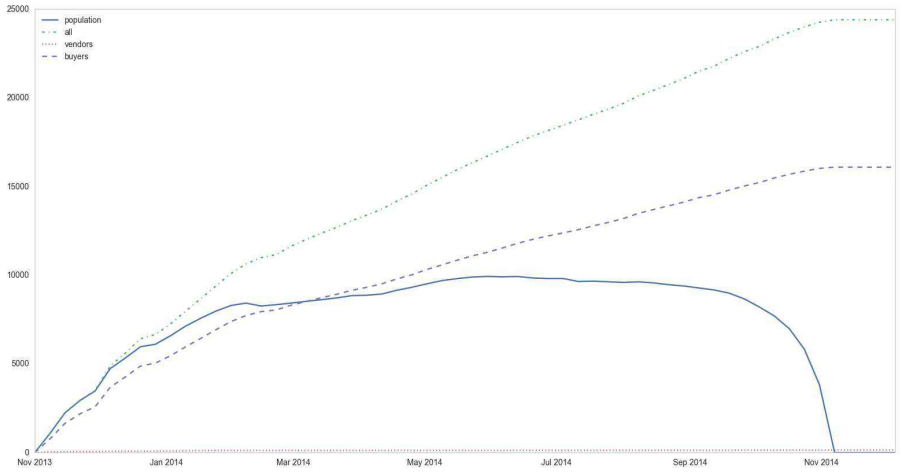

FIGURE 7.7: **Cumulative Population** This plot shows the cumulative population of all users (all), buyers (buyers), and vendors (vendors) against the population (population) over time.

creating accounts, then the change in population each day would be random. The randomness of the population differences from week to week can be measured by comparing this distribution to the Normal Distribution.

The two distributions were measured using the Shapiro–Wilks test (Shapiro and Wilk (1965)) which produced a test statistic of 0.647 and p–value $7.58 \cdot \exp(-27)$. This leads to the conclusion that the data is not normally distributed. If the population fluctuations are not following the Normal Distribution then, this implies that there is something determining the size of certain fluctuations.

Fig. 7.8 shows the size of the population change each week. By inspection, the changes in population are greater at the start and end of the DNM's lifetime. However, other larger than average increases and decreases can be seen around February and July 2014.

To determine which of these fluctuations can be considered out of the ordinary, the day with the largest population fluctuation was systematically removed from the dataset until the results of the Shapiro–Wilks test concluded that the population fluctuations do follow the Normal Distribution.

This process removed 93 data points, 92 of which occurred in either the first or last 3 months of the DNM's lifetime. The majority of data points had to be removed from these early and late periods implying that the population fluctuations occurred at random in the center portion of *Silk Road 2.0's* life but not at the beginning or end.

It can be expected that a DNM would grow rapidly when first launched, or perhaps in waves as new users test out the site and decide to recommend it to others. This behaviour would explain the dramatic population increases described by figure 7.8. However, as *Silk Road 2.0* was closed in Operation Onymous, one would expect that its population would remain constant until the end of its life. This is not observed, instead the population rapidly declined and was potentially already declining before law enforcementclosed the site.

The data point that was removed to make the distribution of fluctuations fit the Normal Distribution and was not in the first or last 3 months of the observation period, corresponded to 14 February 2014. This change was 2 standard deviations from the average change in population (mean 5.25, s.d. 25.1). Here the population fell from 8,389 (on 12 February) to 8,258 (on 14 February) despite the population growing before and after this period. This population fall coincides with the hack on *Silk Road 2.0* in which users collectively lost $2.7 million (DeepDotWeb (2014d)).

The impact of this event lasted until 24 February 2014 when the population finally recovered. In this period, 419 (5%) of the population left the DNM. 242 (58%) of these

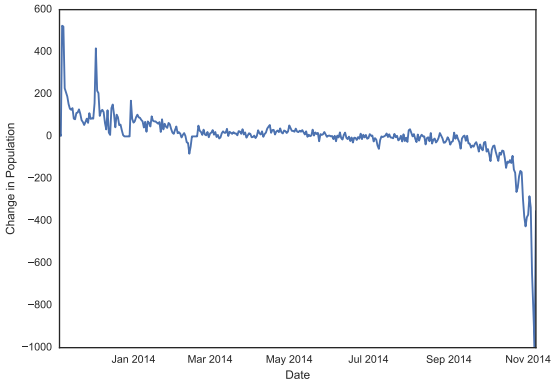

FIGURE 7.8: **Population Fluctuations** This plot shows the change in population each day.

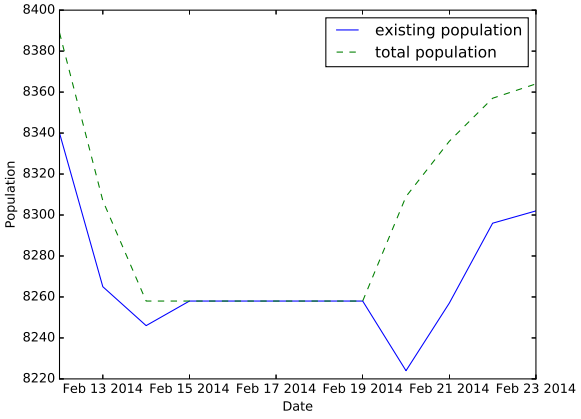

FIGURE 7.9: **Change in Population After Hack** This plot shows change in population and the change in existing population (users who were already on the DNM) immediately after the hack.

users were buyers and 177 (42%) were isolates. The number of sales fell from 628 on 12 February to 329 on 13 February to no sales between 14 February and 19 February which had 291 sales. Similarly, the amount of spending fell from $174,000 to $31,700 to $0 until 19 February when $27,800 was spent.

Figure 7.9 shows the overall population change compared to the change in existing users (users that already had an account on *Silk Road 2.0* prior to the date of the data point). Clearly, existing users reduce the population by leaving.

No other events were found to have a significant impact on the population by this method. This could be because they did not greatly effect the population or because their impact was obscured by the general volatility of the DNM population.

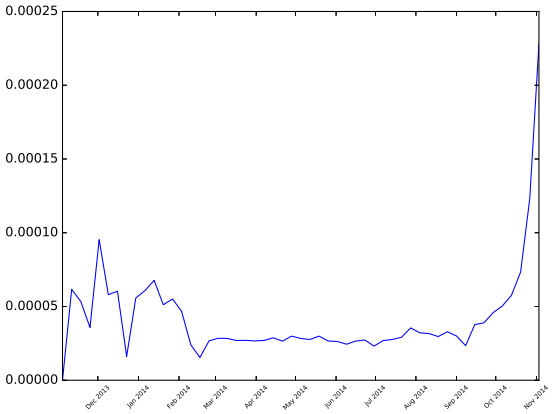

FIGURE 7.10: **Density** This plot shows the density of the network over time.

7.2.2 How Does the Shape of the Network Change Over Time?

The same measures of density and centralisation were applied to each weekly network and the change in each measurement was observed over time. Given that the behaviour of buyers selling back to vendors is neither expected nor observed, reciprocity and transitivity were not also measured over time.

The change in density is displayed in figure 8.4. This shows a sharp decrease in density at the end of 2013 and another decrease in February 2014. These decreases would be expected after events in which the population remains relatively stable but the amount of activity decreases. The decrease in February 2014 implies that the hack of *Silk Road 2.0* had a greater impact on the number of transactions taking place than on the size of the population.

The density then increases just before Operation Onymous when the site closed as the population is rapidly decreasing.

The outdegree (for vendors) and indegree (for buyers) centralisation are plotted over time in figure 7.11. An increase in centralisation shows an increase in dominance for one (or a few) users. The outdegree centralisation is, for the most part, higher than the indegree centralisation and reaches greater maximums however it changes more dramatically. This could be because the extent to which trading fixates on a small number of vendors is greater but, when those vendors lose dominance, the ascension of the new, most popular vendors, is very rapid whereas, for buyers, the most active are not all replaced at once. There are also fewer vendors which means each individual vendor and their activity will have a greater impact on the outdegree centralisation.

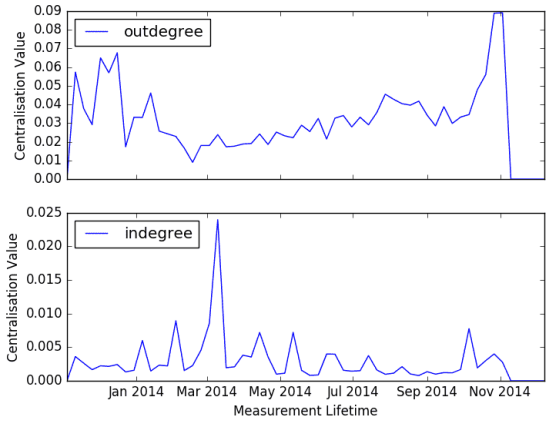

FIGURE 7.11: **Centralisation Over Time** This plot shows the indegree and outdegree centralisation over time.

7.3 Vendor Buyer Relationships

When using publicly available data to understand the impact of law enforcementevents, buyer responses often cannot be measured because they do not have public profiles. Understanding the bond between buyer and vendors can help make assumptions about how this less visible population is affected based on the way that an event makes vendors behave.

First, the characteristics that make vendors popular are explored and described, then the circumstances in which a vendor can cause a buyer to leave the site are evaluated.

7.3.1 What Makes Vendors Popular?

In order to understand if and how buyers clustered around particular vendors, the network was distributed into communities defined by specific vendors.

909 users sold products to 33,123 buyers in the network. Of these, 9 vendors had an exclusive buyer community (i.e. their buyers traded with only them). Each of these vendors had only one transaction. On average, the vendors' buyers traded with an average of 11.6 other vendors (s.d. 6.82, range: 1 60).

380 vendors traded with at least one buyer who traded with them exclusively. 660 vendors traded with buyers such that the buyer made more purchases from that vendor than any other they traded with and, for 504 vendors, at least one of their buyers bought at least half of their items from them.

On average, vendors were responsible for an average of 17.4% of the transactions their buyers made (s.d. 15.1%, range: 0.709 100%). For 129 vendors, at least half of their buyers traded with them the most out of all of the vendors they traded with and, for 58, at least half of their buyers made at least half of their transactions with that vendor.

To understand what made some vendors more popular than others, the ERGM analysis (see Section 5.2.4) described by Duxbury and Haynie (2017) was applied to the network. More precisely, ERGM analysis was used to measure the influence of the following variables: the vendor degree (how many sales they had already made), their reputation (taken to be the cumulative score of their transaction ratings (Duxbury and Haynie (2017))), their total earnings, the affordability (average price) and diversity (number of different types) of their products, and their age (how long their account was active for) on the likelihood of edges being formed in the network.

The statnet package[2] for R 3.4.3 was used to conduct the analysis as a suitable Python library with the same capabilities could not be sourced.

The network was too large to conduct analysis on in its entirety. Instead, the analysis was conducted on each weekly snapshot and averaged. Not all of the variables were significant for each week and the number of weeks each variable was significant is given in Table. 7.3.

The significant results were aggregated for each of the variables that were significant for most weeks (i.e., all the variables except for the non-EU locations). These values are presented in Table. 7.4.

The baseline probability of a tie being formed between a buyer and a vendor ranged between $6.81 \cdot \exp(-7)$ (recorded on 28 April 2014) and 0.000611 (recorded on 11 November 2013), i.e., in the model any edge of value \$1 had a $6.81 \cdot \exp(-5)\%$ chance of being formed from a buyer to a vendor when this probability was at its lowest and a 0.0611% chance at its highest.

The different variables either increased or decreased this probability. For example, an increase in the age of a vendor nearly always increases the probability that an edge will be made to that vendor (it only decreased the probability in one instance, on 10 February 2014). At its highest impact, the probability increases from $3.717 \cdot \exp(-3)\%$ to $3.720 \cdot \exp(-3)\%$, i.e. a one unit increase in a vendor's age increases the probability an edge will be made to them by a factor of 1.001. The diversity of the products sold by the vendor had a similar, positive impact. At its highest impact, a one unit increase in reputation increases the probability an edge will be formed by a factor of 1.03.

[2] http://www.statnet.org/

TABLE 7.2: Influence of Vendor Characteristics on the Probability an Edge is Formed

	Probability of an Edge Being Formed / Baseline Probability		
Variable	Maximum	Minimum	Average
Europe	2.38	0.111	0.576
Reputation	1.03	0.999	1.00
Age	1.01	0.999	1.00
Diversity	1.01	0.997	1.00
Affordability	1.00	1.00	1.00
Earnings	1.00	1.00	1.00
Vendor Degree	1.00	0.995	1.00

The variable with the most dramatic impact on the probability of a tie being created was whether or not the vendor was European. None of the other nationalities were significant enough to be included in the results (presumably either because too many vendors held this nationality (for the UK or unknown) or not enough (for the remainder of the nationalities)). In 50 of the 52 observed weeks, being European decreased the probability that a buyer would purchase from them however on 20 January and 17 February 2014, being European increased the probability of a buyer making a purchase from them 2.38 and 2.33 times respectively.

The remaining variables did very little to effect the probability of a tie being formed at most increasing or decreasing it by a few hundredths of a percent. The impact of each of the variables on the probability of a tie being formed is given in table 7.2.

Figure 7.12 shows the probability of a tie being formed due to each variable being increased by 1 unit over time. This shows that the probability of edges being formed in the network does not remain constant over time, instead it decreases after the establishment of the site and falls again after the hack of *Silk Road 2.0*.

It can be seen that the influence of most variables (all but location) follows the same pattern as the base probability. This implies that there were potentially events that influenced the likelihood that buyers made purchases but these did so universally and not in a manner that favoured particular vendors.

The results of the ERGM analysis differ from those presented by Duxbury and Haynie (2017) who found that vendor reputation had the largest influence on the probability that a transaction would be made. They found that this characteristic had nearly 10 times the impact as it did in this study. This is despite the fact that vendor reputations varied more greatly in our dataset.

The greater influence could have been caused by the fact that the dataset used by Duxbury and Haynie (2017) is smaller, with both the indegree of vendors and outdegree of buyers being much smaller than in the dataset used in this study, which may have

TABLE 7.3: Number of Weeks Each Variable Held Each Level of Significance

Variable	None	$p < 0.1$	$p < 0.05$	$p < 0.01$	$p < 0.001$
Edges	0	0	0	0	52
Vendor Degree	0	0	0	0	52
Trustworthiness	7	1	0	2	42
Earnings	3	1	1	2	45
Affordability	1	0	0	1	50
Diversity	2	0	0	1	49
Age	2	0	1	1	48
UK	52	0	0	0	0
America	42	3	2	3	2
Africa	52	0	0	0	0
Asia	52	0	0	0	0
Europe	0	0	0	0	52
Oceania	50	1	1	0	0

TABLE 7.4: Measures of Variation for the Statistically Significant Values of the Results of ERGM Analysis Over Time

Variable	Mean	Standard Deviation	Maximum Value	Minimum Value
Edges	-11.9	1.75	-7.4	-14.2
Vendor Degree	0.000623	0.00155	0.00479	-0.00466
Trustworthiness	0.00159	0.00581	0.0337	-0.00134
Earnings	$7.79 \cdot \exp(-5)$	0.000296	0.00178	$-1.69 \cdot \exp(-5)$
Affordability	0.106	0.0497	0.181	0.0214
Diversity	0.00315	0.00180	0.00739	0.00299
Age	0.00315	0.00135	0.00595	-0.000970
Europe	-0.661	0.432	0.869	-2.20

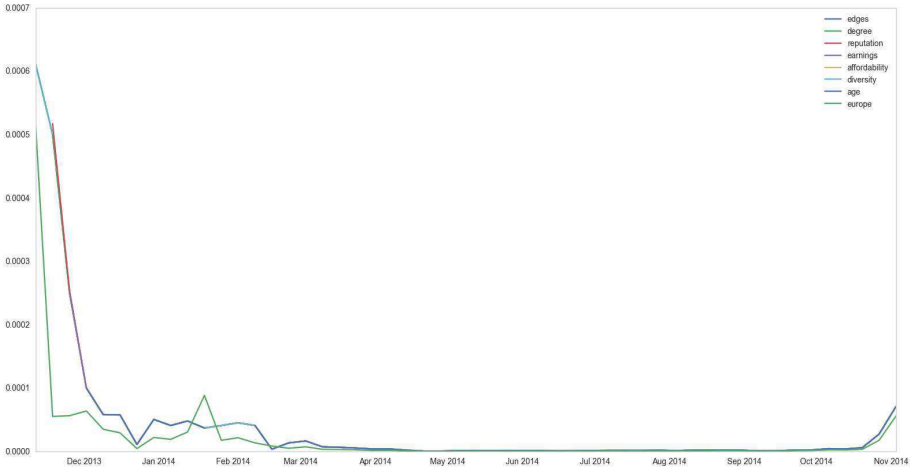

FIGURE 7.12: **Probability of an edge** being formed given a one unit increase in each of the variables.

concentrated the influences of certain vendor characteristics. Alternatively, as Duxbury and Haynie (2017) restrict their analysis to the opioid network, the result may reflect that opioid buyers are more discerning in their purchases and/or opioid vendors provide a better service.

7.3.2 Can Vendor Behaviour Influence Buyers to Leave the Market?

In order to understand the influence that vendors had on buyers and determine the extent to which buyers were participating in the system because of particular vendors, the impact that vendors had on their buyers by leaving was measured. For 24 vendors, none of their buyers maintained their accounts after the vendor left the site. For 241 vendors, none of their buyers maintained their accounts for longer than a week after the vendor left the site.

On average, 32.8% of each vendor's buyers closed their accounts when they did (s.d. 34.9%) and 42.0% of their buyers closed their account within a week (s.d. 41.1%). For 288 vendors, all of their buyers maintained their accounts after they left the site.

For 62 vendors, none of their buyers made a purchase after they left the site. For 282 vendors, none of their buyers made purchases for longer than a week after they left the site. On average, 47.1% of a vendor's buyers would stop trading after they left the site (s.d. 36.6%) and 54.0% would stop trading within a week after the vendor left (s.d. 39.0%). For 220 vendors, none of their buyers stopped trading when they left and for 192, none of their buyers stopped trading within a week after they left.

Logistic Regression was used to see which types of buyers continued to trade and which stopped trading after a vendor left. The variables included in the model were designed to capture how experienced the buyer was, how much they valued the vendor, the experience of the vendor, the activity of the buyer after the vendor leaves, and their final transaction experience.

The experience of the buyer was measured using the age of their account, the number of purchases they had made, the number of different types of products they had bought, the amount they spent, the number of vendors they had bought from and the number of vendors they had previously traded with who had already left the site.

The value they placed on the vendor was measured using the comparative rating they had given them (the average rating given to that vendor minus the average rating they had given across all their transactions), the proportion of purchases made from that vendor, the comparative length of their relationship, the comparative proportion of transactions that were cancelled, and the proportion of their total spending spent with that vendor.

The vendor experience was measured by the vendor's age, the number of buyers they have sold to, the vendor reputation, and the total the amount they earned by the vendor.

The activity of the buyer in the week after the vendor leaves was measured by the number of new transactions they made and the number of new vendors they traded with.

Finally, the final transaction experience of the buyer was measured by the most recent rating they had given for any purchase and the most recent rating they had given to the vendor.

To conduct the analysis the libraries Pandas and SKLearn for Python 2.7 were employed[3].

The model was first applied to buyers who closed their accounts immediately after a vendor they were trading with left, then to buyers who closed their account within a week of the vendor leaving, then to buyers who stopped making purchases immediately after a vendor left and finally to buyers who stopped making purchases within a week after a vendor they were trading with left.

The models that evaluated buyer behaviour in the week after a vendor leaving are more accurate than those that looked at buyer behaviour straight afterwards, implying that buyer reactions are not immediate. Further, the models that looked at if a buyer stopped trading, as opposed to shutting down their account, are also more accurate, implying that the buyer behaviour may be passive rather than active.

The results of the model that tried to predict which buyers stopped trading within a week of the vendor leaving are now discussed. There were 51,598 buyer - vendor pairings from instances of vendors leaving to evaluate. In 43,032 (83%) instances, the buyer stopped trading within a week of the vendor leaving and in 8,566 (17%) instances they did not.

First, the variables were inspected to see if they had a significant impact on the model. The significance of each variable is given in Table. 7.5. This was not the case for the comparative length of the buyer/vendor relationship and the proportion of total spendings that the buyer spent with the vendor. And, as such, they were removed from the model.

The model has an accuracy of 0.837. To improved its accuracy, Recursive Feature Elimination (RFE) was used to systematically remove variables from the model and evaluate if this action increased its accuracy. By removing the total amount spent by the buyer and the total profit made by the vendor the accuracy of the model was increased to 0.893.

[3] https://towardsdatascience.com/building-a-logistic-regression-in-python-step-by-step-becd4d56c9c8

TABLE 7.5: Logit Regression Results for the Model to Predict if a Buyer will Stop Trading when a Vendor Leaves the Market

	Coefficient	Standard Error	z	$P > \|z\|$
Age of Buyer	0.0045	0.000	21.626	0.000
Number of Purchases	0.0149	0.002	7.507	0.000
Diversity of Purchases	-0.0730	0.006	-12.426	0.000
Amount Spent	$-1.341 \cdot \exp^{-5}$	$3.06 \cdot \exp^{-6}$	-4.380	0.000
Number of Vendors Purchased From	0.3415	0.008	42.437	0.000
Number of Vendors Buyer has Traded With that have Left	-0.6429	0.011	-56.119	0.000
Comparative Rating Given to Vendor	0.0877	0.013	6.634	0.000
Proportion of Purchases Made With Vendor	0.2088	0.143	1.464	0.143
Comparative Length of Relationship	-0.0004	0.000	-0.745	0.456
Comparative Proportion of Cancelled Transactions	1.2645	0.080	15.748	0.000
Proportion of Total Spending Spent on Vendor	-0.1606	0.128	-1.250	0.211
Age of Vendor	-0.0150	0.000	-64.942	0.000
Number of Buyers Vendor has Traded With	0.0009	$7.18 \cdot \exp^{-5}$	12.116	0.000
Vendor Reputation	-0.0001	$1.08 \cdot \exp^{-5}$	-13.469	0.000
Total Amount Earned by Vendor	$1.484 \cdot \exp^{-6}$	$1.46 \cdot \exp^{-7}$	10.181	0.000
Number of New Transactions	0.4608	0.033	13.956	0.000
Number of New Vendors Traded With	0.1895	0.057	3.354	0.001
Last Rating Given	0.2454	0.009	27.845	0.000
Last Rating Given to Vendor	-0.0640	0.011	-5.573	0.001

TABLE 7.6: Accuracy of the Model to Predict if a Buyer will Stop Trading when a Vendor Leaves the Market

	Precision	Recall	F1-Score
Buyers who Maintained their Account	0.73	0.56	0.64
Buyers who Closed their Account	0.92	0.96	0.94
Avg / total	0.89	0.89	0.89

10-fold cross validation was used to check the consistency of the model, the mean of the results of this was 0.891. The classification report of the final model is given in Table. 7.6.

The multicollinearity of the variables in the model was measured by calculating their VIF scores. These are given in 7.7.

The buyer characteristic that most affected if a buyer would continue to trade was the

TABLE 7.7: The Variance Inflation Factor of Each Variable in the Logistic Regression Model.

Variable	VIF Score
Age of Buyer	1.6
Number of Purchases	5.0
Diversity of Purchases	10.0
Amount Spent	1.2
Number of Vendors Purchased From	10.6
Number of Vendors Buyer has Traded With that have Left	4.8
Comparative Rating Given to Vendor	2.2
Proportion of Purchases Made with Vendor	6.1
Comparative Length of Relationship	1.2
Comparative Proportion Cancelled Transactions	1.3
Proportion of Total Spending Spent on Vendor	5.4
Age of Vendor	1.4
Number of Buyers Vendor has Traded With	56.1
Vendor Reputation	68.8
Total Amount Earned by Vendor	9.1
Number of New Transactions	1.4
Number of New Vendors Traded with	1.4
Last Rating Given	1.3
Last Rating Given to Vendor	2.6

number of trading partners they had who had left – the fewer trading partners they had experienced leaving, the more likely they were to continue trading. To a lesser degree, the older a buyer was and the more purchases they had made, the more likely they were to continue trading. The amount the buyer had spent was a very marginal negative influence on them continuing to trade. This amounts to buyers being more likely to continue trading if they are more experienced and have less experience of vendors leaving.

The variables that capture the buyer's behaviour after their trading partner leaves both positively influence the model. This means that the more purchases and, to a lesser degree, the more trading partners, a buyer had in the week after the vendor left the more likely they were to continue trading. This is presumably because they were able to find options to replace the vendor they had been trading with. Alternatively, this variable could indicate that the buyers more likely to continue trading were those with a more regular purchasing habit.

The amount that the vendor characteristics influenced the buyer's decision to continue trading was small. The younger the vendor was, the more likely the buyer would continue to trade without them. However, the characteristics of the vendor appear to have had a much lower influence over the buyer's decision.

Only three characteristics that measured the buyer/vendor relationship were significant and only two of these had low enough VIF scores to be considered in the model. These were how the buyer rated the vendor, relative their other trading partners, and the last

rating the buyer gave the vendor. The comparative rating the buyer gave the vendor was a positive influence on if they continued to trade or not, i.e. buyers who rated the vendor higher than other vendors they traded with were more likely to continue trading. If buyers had a cancelled a disproportionately high number of transactions with the vendor they were more likely to continue trading. This variable had the largest influence on the model and potentially represents that buyers who either avoided being scammed when the vendor left *Silk Road 2.0* or who had a poor experience of the vendor were more likely to continue trading.

The lower the last rating the buyer gave the vendor the more likely they were to continue trading which supports the idea that buyers who had a poor experience trading with a vendor continued to trade. The last rating that the buyer gave before the vendor left had a much bigger influence on the model and positively influenced the model, however. It may, therefore, have not been enough for buyers to have a poor last transaction with the vendor but they may also have had to have experienced positive transactions with other trading partners.

This analysis implies that, as opposed to buyers stopping trading because a particularly favoured vendor has left the DNM, buyer decisions are more likely to be affected by their previous buying experiences and current buying options. If buyers had experienced multiple trading partners leaving then potentially they felt that it was not worth staying on *Silk Road 2.0*, similarly if they could not find a suitable replacement to trade with they may have looked elsewhere. As the buyer's perception of their last purchase appears to have been a contributing factor to their decision, the way in which bad transaction experiences may motivate buyers to stop trading is explored in greater depth.

7.3.3 Bad Transaction Experiences

A bad transaction experience was defined as any transaction which received a rating of 1 or 2 out of 5. This metric was chosen because it is defined by the user and, therefore, does not require imposing a subjective measure of success onto the data. It also allows for transactions that contain little review data to be included in the analysis.

5,904 of the transactions were negative and 3,766 users (who have accounts in the database) experienced a negative purchase. The number of transactions assigned each rating is given in Table. 7.8.

112 (79%) of UK vendors received at least one negative review. Figure 7.13 shows the proportion of negative reviews a vendor has received plotted against their lifetime. This

TABLE 7.8: The Distribution of Ratings Across All Reviews

Rating	Number of Reviews	Frequency
1	5,288	2.78%
2	616	0.323%
3	1,545	0.810%
4	3,327	1.74%
5	180,026	94.4%

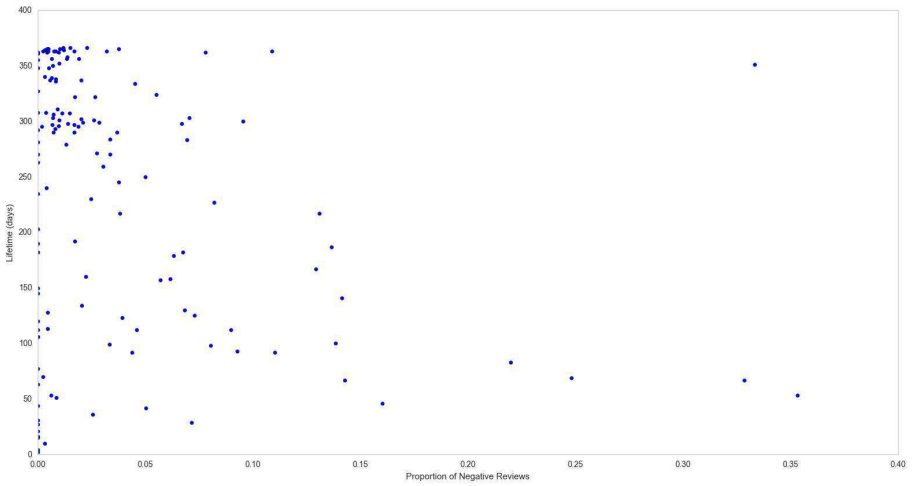

FIGURE 7.13: **Review Quality and Lifetime** This plot shows relationship between the proportion of negative reviews a vendor receives and the length of their lifetime.

shows that, whilst some of the vendors with the highest proportion of negative reviews had shorter lifetimes, the two attributes are not clearly negatively correlated.

However, if a buyer had a bad transaction experience, they may have chosen to stop trading on the site because they were concerned about losing more money or purchasing low quality product.

Figure 7.14 shows the proportion of negative of reviews a buyer gives plotted against their lifetime. This does not show the expected strong negative correlation between proportion of negative reviews and lifetime.

Other factors that may have contributed to this decision include the buyer experience, measured by their age, the number of transactions they had already made and the number of previous bad experiences they had already had. Additionally, the cost of the product may have been a factor as buyers who lost more money may have been more unwilling to continue trading. Finally the relationship between the buyer and the vendor could be an influence - if the buyer had traded with the vendor before they may be more willing to forgive a poor transaction experience, as they might also with specific well

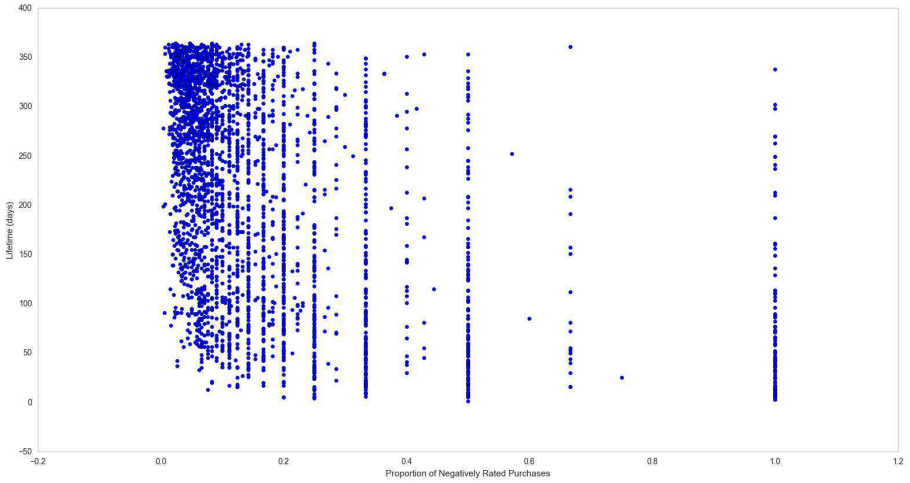

FIGURE 7.14: **Transaction Quality and Lifetime** This plot shows relationship between the proportion of negative reviews a buyer gives and the length of their lifetime.

TABLE 7.9: Accuracy of the Model

	Precision	Recall	F1-Score
Stopped Trading	0.30	0.00	0.00
Continued Trading	0.93	1.00	0.96
avg / total	0.87	0.93	0.90

known vendors or vendors with high reputations, as such all these variables were also collected.

Logistic Regression Analysis was conducted on 1,940,540 transactions to evaluate the influence the previously described variables had on whether a buyer chose to continue trading or not. In 1,808,674 of the instances they did and, in 131,866 of the instances they did not make any further purchases.

The logistic regression results are presented in Table. 7.10 which shows that all of the variables were significant, except the rating of the vendor.

They show that the biggest contributing factors to a buyer continuing to trade appear to be the rating they gave their transaction and if they had traded with the vendor before. As expected, the higher the transaction rating, the more likely they were to continue trading. If they had traded with the vendor before they were more likely to make another purchase, this could be because they were less put off by bad transactions if they trusted the vendor they were trading with and thought they could be compensated.

TABLE 7.10: Logit Regression Results of the Model to Predict if a Buyer will Stop Trading after a Bad Transaction Experience

	Coefficient	Standard Error	z	$P > \|z\|$
Age of Buyer	0.0043	$2.96 \cdot \exp(-5)$	143.754	0.000
Transaction Rating	0.2314	0.002	102.186	0.000
Number of Purchases	0.0078	$8.45 \cdot \exp(-5)$	91.931	0.000
Number of Bad Reviews	0.1191	0.003	41.471	0.000
Cost	$-2.994 \cdot \exp(-5)$	$1.71 \cdot \exp(-6)$	-17.498	0.000
Vendor Reputation	$-2.089 \cdot \exp(-5)$	$2.01 \cdot \exp(-7)$	-103.886	0.000
Have they Previously traded?	0.3180	0.010	30.450	0.000

To a lesser degree, the age of the buyer and the number of purchases they had already made positively influenced the model. This implies that more experienced buyers are marginally more likely to continue trading after bad transaction experiences.

The reputation of the vendor and the amount the buyer paid in the transaction each had a very marginal influence on the model. The more expensive the transaction was, the less likely the buyer would continue to trade potentially because, in the instances where the transaction was bad the buyer had lost more and this could make them more risk adverse in the future. If the reputation of the vendor was higher, the buyer was marginally more likely to continue trading.

A surprising result was that the number of bad reviews the buyer had previously given had a positive influence on if they traded or not. Potentially some buyers gave a high number of negative reviews over a long lifetime and this skewed the data, indeed the maximum number of negative reviews given by a buyer who continued to trade was 163 compared to 46 for buyers who stopped trading. Figure 7.13 implies that the number of negative reviews is not strongly correlated with a buyer's lifetime supporting the argument that the quantity of negative reviews given does not necessarily influence a buyer's decision to continue trading.

The model had an accuracy of 0.932, this was checked using 10-fold cross validation which returned a score of 0.932 also. The full classification report is given in Table. 7.9. This report implies that the large skew towards continuing to trade aided in the accuracy of the model as it could be guessed, with high probability, that buyers would continue to trade.

The VIF scores of each variable is given in table 7.11[4].

[4]These values were calculated by averaging the values of multiple random samples as the dataset was took big to evaluate in its entirety. The VIF scores of each random sample were extremely similar and differed by, at most, a value of .1.

TABLE 7.11: The Variance Inflation Factor of Each Variable in the Logistic Regression Model.

Variable	VIF Score
Age of Buyer	1.0
Transaction Rating	1.6
Number of Purchases	1.0
Number of Bad Reviews	1.6
Cost	1.0
Vendor Reputation	1.0
Have they Previously Traded?	1.0

TABLE 7.12: The Variance Inflation Factor of Each Variable in the Logistic Regression Model.

Variable	VIF Score
Age of Buyer	1.0
Transaction Rating	1.6
Number of Purchases	1.0
Number of Bad Reviews	1.6
Cost	1.0
Vendor Reputation	1.0
Have they Previously Traded?	1.0

Whether or not a buyer continues to trade was determined simply by if they made another purchase. However some buyers make multiple purchases on the same day which is an insufficient time for them to react to a bad purchase. Therefore, this analysis was repeated where buyers were considered to be continuing to trade if they made another purchase from the day after the transaction being examined took place.

With this new definition, for 263,401 transactions the buyer stopped trading and for 1,677,139 the buyer did not.

Table 7.13 shows that this adaptation to data classification does not have a large impact on how the variables influence the model. However, the model was less accurate in this instance as it only had an accuracy of 0.863 (10-fold cross validation produced a mean of 0.862). Once again, the accuracy of the model is very close to the proportion of instances where the buyer continues to trade indicating that the model is not utilising the variables to categorise the transactions and, instead, is able to guess with high accuracy because of the sample skew.

The VIF scores of each variable is given in table 7.12.

TABLE 7.13: Logit Regression Results of the Model to Predict if a Buyer will Stop Trading 1 Day after a Bad Transaction Experience

	Coefficient	Standard Error	z	$P > \|z\|$
Age of Buyer	0.0059	$2.25 \cdot \exp(-5)$	263.592	0.000
Transaction Rating	0.0866	0.002	45.283	0.000
Number of Purchases	0.0012	$3.1 \cdot \exp(-5)$	39.373	0.000
Number of Bad Reviews	0.0752	0.002	39.811	0.000
Cost	-0.0002	$5.35 \cdot \exp(-6)$	-29.900	0.000
Vendor Reputation	$-1.968 \cdot \exp(-5)$	$1.62 \cdot \exp(-7)$	-121.604	0.000
Have they Previously traded?	0.0593	0.009	6.616	0.000

Chapter 8

Reddit Forums Results

This section presents the results of the analysis on the subreddits */r/DarkNetMarkets*, */r/dnmuk*, */r/Ebay*. It begins by describing the general statistics of each forum and compares the size, shape, and proportion of deleted content of the DNM related subreddits to */r/Ebay* and then evaluates the impact of Operations Hyperion and Bayonet.

The results presented in this chapter are used to evaluate the hypotheses 5 and 17 which describe how an adverse event might impact a user if they perceive the ecosystem to be less stable or themselves to more at risk from law enforcement. They are also used to evaluate the following hypotheses: 8, 9, 10 and 18 which describe the potential affects of losing an account or having a heightened concern of the presence of law enforcement. Finally, these results are used to evaluate Hypothesis 4 which describes how a reduction in convenience affects the resilience of the ecosystem.

8.1 General Statistics

In this section, each subreddit is described in terms of its user statistics and network characteristics. The characteristics of the subreddits are then compared to identify differences that may be attributed their content.

The size of each subreddit is measured by counting the number of contributor accounts and contributions, in the form of posts and comments. Additionally, these values are used to measure the proportion of material which is deleted. They show that the subreddit */r/DarkNetMarkets* is substantially larger than */r/dnmuk* and that both DNM related subreddits have had more material deleted than */r/Ebay*.

The amount of engagement is measured using the length of time that posts, comments and accounts are active for and the amount of posts/comments per account. These measurements are used to understand if contributors engage with the same conversations for long periods of time or if these topics change rapidly. Similarly, they are used to understand if contributors are maintaining the same account across multiple conversations.

These statistics show that, in each subreddit, posts received very few comments and were not active for very long (on average, less than a day). Contributor accounts made very few posts and comments and, on average, only lasted a few days. These statistics were similar for each of the subreddits.

Finally, each subreddit is described by the characteristics of its network by measuring its density, transitivity, and centrality. These measurements are used to understand the connectivity of each network which is shown to be sparse for both DNM subreddits and marginally less sparse for */r/Ebay*. None of the networks had high centralisation values, i.e. the conversations were spread across a large number of posts, rather than being centred around a small number of highly popular posts.

These statistics are used to conclude that users of DNM subreddits actively engage (by making posts and/or comments) on a short term basis either because contributors only need to engage sporadically or because they deliberately change accounts after each engagement. Though there are some posts that receive a lot of activity, each only receives a small fraction of the total amount of activity and most receive only 1 comment.

The comparison between these subreddits and */r/Ebay* implies that it is unlikely this behaviour is driven by the content of the subreddits, though potentially DNM contributors delete their accounts more readily than E-bay contributors.

8.1.1 /r/DarkNetMarkets

As discussed in Chapter 4, the subreddit */r/DarkNetMarkets* contained 324,120 posts and 572,585 comments. On average, each post had 1.77 comments (with a minimum of 0 and maximum of 895). The number of comments per post is dramatically positively skewed with 98% of posts receiving less than 10 comments and 71% of posts receiving only one comment. Despite this, only 625 posts (<1%) did not receive any comments. Just 13 posts received more than 100 comments and, whilst much of the content around these posts was removed, they appear to cover topics from Bitcoin, or other product, give-aways to issues with Coinbase, to life advice.

The lifetime of each post was measured as the number of days between the post first being posted and the date of the last comment. On average, posts were active for just

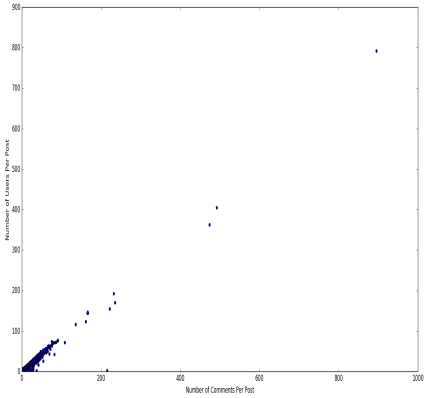

FIGURE 8.1: **Number of Users vs Number of Comments** this figure shows the number of comments and number of users for each post.

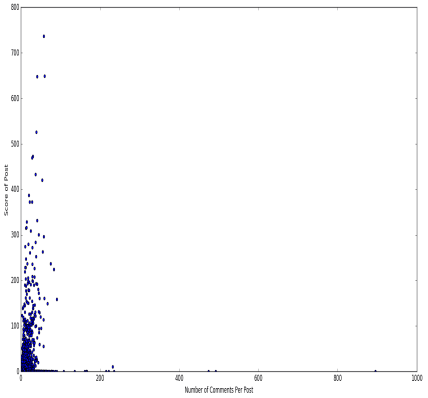

FIGURE 8.2: **Score vs Number of Comments** this figure shows the score of each post plotted against the number of users.

less than one day with 95% of posts remaining active for less than a day. 4,585 posts (1%) were active for more than one week and these posts had, on average 5.35 comments (max 474, min 1) which means that the two posts with the most comments were active for less than a week. The post that was active for the longest period of time was active for 30 days and received only 2 comments.

Each posts had, on average 2.71 contributors, with a maximum of 792 and minimum of 1. The majority (72%) of posts had only 2 contributors which comprised of 1 poster and 1 commenter in almost all instances. Figure 8.1 shows that, in most instances, the number of comments is proportional to the number of users.

Posts were each assigned a score calculated from the amount of up and down votes they received. Posts overwhelmingly (97%) received a score of 0. The average score was 0.208 with a maximum of 736. Figure 8.2 shows that there is a trend of positive correlation however many highly commented posts did not receive any score.

There were 417,701 contributors to the forum during the measurement period. Of these, 103,227 (25%) made at least one comment. The average number of comments was 1.37, with 82,479 (80%) of commenters making one comment only. The maximum number of comments made by one user was 20,040 and the user was 'AutoModerator', the next highest number of comments was 5,512 made by an account with a name that did not imply it was a bot or admin account.

320,465 (77%) contributors made at least one post. The average number of posts per contributor was 0.776 with the 'AutoModerator' making the most posts at 139. The

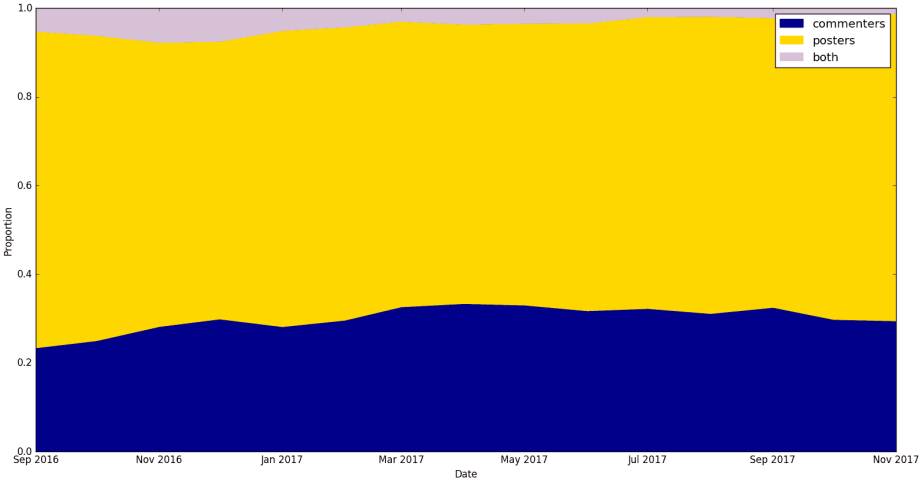

FIGURE 8.3: **Proportion of Contributor Type** this plot shows the proportion of contributors who made posts, comments or both during the measurement period.

next highest number of posts made by anyone user was 54. 99.6% of contributors who posted, posted once.

5,991 (<1%) contributors made both at least one post and one comment. They made, on average 1.59 posts and 33.1 comments each.

Contributor accounts lasted on average 5.22 days with 4 accounts being present for the entire measurement period and 400,017 (96%) lasting less than 1 day. The average account length for commenters was 18 days and 1.4 days for posters. All 4 of the contributors who were present for the whole measurement period made both posts and comments, this group had the highest average account length - 134 days. The proportion of contributors who made posts, comments, and posts and comments is plotted in figure 8.3. It shows that, whilst the proportion of users who post and those who comment remain relatively stable, the proportion of users who both comment and post diminishes over time.

These statistics show that, for most users, when they actively engage with the forum (i.e. make a contribution) they either do so rarely, use different accounts to make multiple contributions, or delete their contributions. This supports an understanding of the user population that either chooses to use the forum sparingly (for instance when they have a specific question) and/or who is cautious about the information they share. Analysis of the forums cannot shed light on how frequently users passively engaged with the forums (by reading content) or actively engaged through private channels.

8.1.1.1 Network Analysis

A network was created with nodes representing users and edges demonstrating one node had commented on a post of the other. The density of the network is $1.73 \cdot \exp(-5)$, i.e. the network is extremely sparse[1]. This is likely due to the large proportion of users who only interacted with the forum once but also implies that few contributors commented on multiple posts. However, the average number of posts commented on by commenters who made more than 1 comment was 23 and only 130 (<1%) of these commenters commented on just 1 post.

The reciprocity of the network is 0.00122. This measures the extent to which a poster posts on a post made by one of the contributors who commented on their post. Whilst this value is low, a more useful measure here is the number of "conversations" that occurred in the forum, i.e. the number of times that a person replied to a comment and then had their reply replied to. To measure this, a new graph would need to be created where an edge is drawn from one user to another if they replied to their comment. However, the data source did not contain fields that showed when a comment was a reply to another comment or embedded someone else's comment in their own. Instead it was assumed that, for each comment, the comment that occurred chronologically afterwards was a direct reply to it. These new edges were added to the graph and reciprocity was recalculated as 0.0280. Therefore the true reciprocity is between 0.00122 and 0.0280.

The transitivity of a network measures the number of triangles found in the network, i.e. if a talks to b who talks to c, in how many instances does c also talk to a. This value was 0.000329 and rose to 0.00561 when edges were added to represent the comments as responses to other comments.

The centralisation measures how different the network is from a star network, i.e. a network focused around one central node who is connected to everyone else. The centralisation of the network is 0.0515, the indegree centralisation is 0.0477 and the outdegree centralisation is 0.00379. This means that posters are more central than commenters (because edges are drawn from post to comment and the indegree centralisation is larger than the outdegree centralisation).

Over time, these values changed and so 15 separate graphs were made across monthly intervals. The measurements were plotted in figure 8.4, 8.5, 8.6, 8.7.

[1]The density of the network would be lower if comments from bots, such as the account 'AutoModerator' were to be removed.

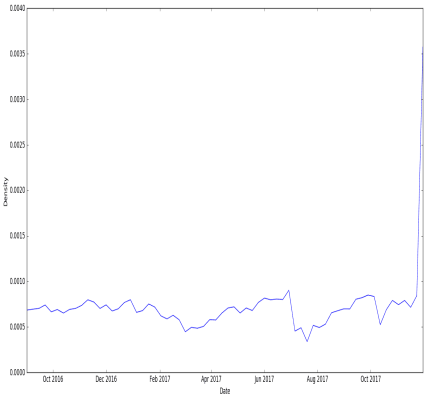

FIGURE 8.4: **Density** this figure shows the change in network density over time for subreddit */r/DarkNetMarkets*.

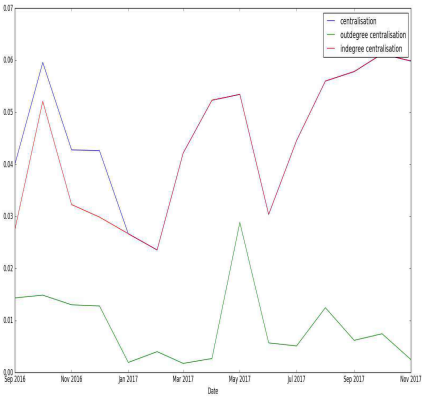

FIGURE 8.5: **Centralisation** this figure shows the change in outdegree, indegree, and overall centralisation over time for subreddit */r/DarkNetMarkets*.

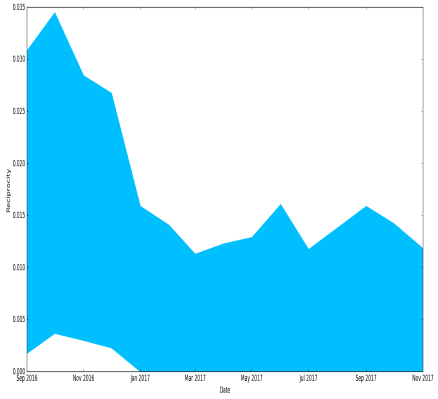

FIGURE 8.6: **Reciprocity** this figure shows the change in network reciprocity over time for subreddit */r/DarkNetMarkets* as a range between the measure on the network formed with the minimum and maximum number of possible edges.

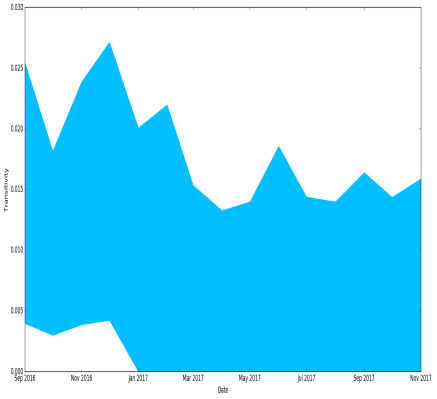

FIGURE 8.7: **Centralisation** this figure shows the change in transitivity over time for subreddit */r/DarkNetMarkets* as a range between the measure on the network formed with the minimum and maximum number of possible edges.

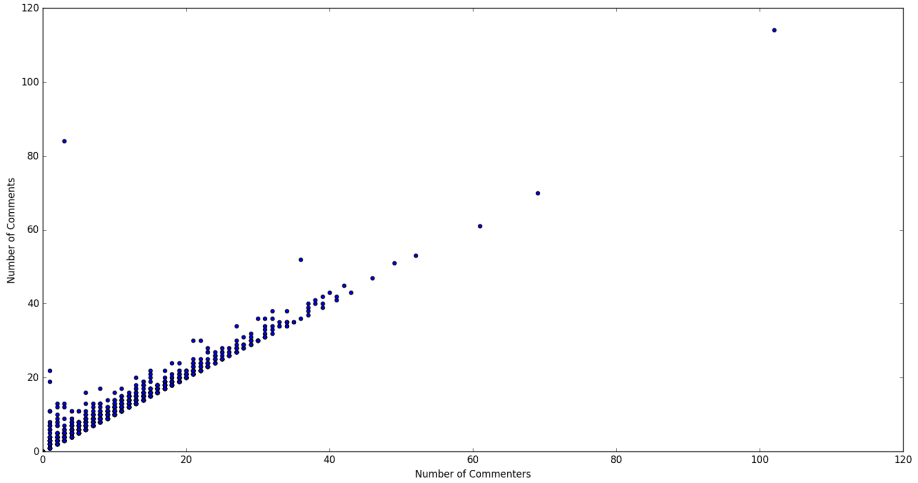

FIGURE 8.8: **Number of Users vs Number of Comments** this figure shows the number of comments and number of users for each post.

8.1.2 /r/dnmuk

During the measurement period, 168,873 posts and 281,248 comments were recorded. There were 1.67 comment per post and the post that received the most comments received 114 comments and appears to be a post about Brexit posted on 5 January, 2017. 204 (<1%) posts received no comments at all.

In addition, posts had an average of 1.63 different posters and figure 8.8 shows a strong positive correlation between the number of contributors to a specific post and the number of comments it received implying that, for most posts, contributors only commented once.

Most posts (97%) were active for less than a day and the average number of days a post was active for was 0.0839. The longest any post lasted was 28 days, there were two posts active for this length of time, one was an 11 comment discussion about a vendor which transitioned from recommendations to a discussion of missing packages and the other only had two comments and was about the ability to buy over the counter (OTC) in New York.

There were a total of 200,792 contributors in the dataset. Of those contributors, 167,919 (83%) left at least one post, 33,852 (17%) left at least one comment, and 979 (<1%) left at least one comment and at least one post.

The most posts made by any one contributor was 24 (username 'mastersdrummer') and, on average, contributors made 0.841 posts each with the vast majority of contributors

(99.9%) making less than 5 posts. The most comments made by any one contributor was 6,925 made by a contributor named 'GroovyEFS' who also made 6 posts. On average, contributors who commented made 8.31 comments, however the number of comments made by contributors is still negatively skewed with 94% of commenting contributors making less than 10 comments.

On average, contributors actively engaged with the forum for 2 days and 98% of contributors actively participated for less than 1 day. The longest active lifetime for any contributor was 455 days (the length of the measurement period) and 3 contributors ('thenorm123', 'crakbenz', and 'YoMommaRollsMyWeed') were active for this period.

The density of the network constructed from posters and commenters was $4.95 \cdot \exp(-5)$. The reciprocity was 0.00141 and the transitivity 0.000368. These values are very similar to */r/DarkNetMarkets* and, indeed, many of these statistics are similar across both forums. It is likely, therefore, that users engaged with them in a similar manner.

8.1.3 /r/Ebay

/r/Ebay contained 54,606 posts, 87,264 comments and 64,428 contributors, 5,798 of which had the username '[deleted]'. 52,329 (81%) of the contributors made a post. On average, contributors made 0.848 posts and the maximum number of posts made by one user was 1,606 (made from an account whose username did not imply it was admin). Only 5 users made more than 10 posts.

14,012 (22%) of contributors made at least one comment. On average, contributors made 1.35 comments and the maximum number of comments made by one user was 3,657. 1,025 users made more than 10 comments. Just 1,913 (3%) of users made both a comment and a post.

Contributor accounts lasted for, on average 7.33 days. 7 accounts were present throughout the whole measurement period and those accounts each made at least 100 comments (making 992 each on average) though only 5 made a post and the maximum number of posts made by one contributor was 8 (average 3.14). 95% accounts lasted for less than a day.

The posts had, on average, 1.60 comments in response to them. The maximum number of comments received by anyone post was 25, two posts had 25 comments, one post was about the variability of sales and if business was particularly bad at the time of posting whilst the other was seeking advice on potential customers who try to haggle on the price of an item. 2,120 (4%) of posts had no comments and 99% had less than 10 comments.

Posts lasted, on average, 0.0712 days and 52,744 (97%) posts lasted less than a day. The post that remained active for the longest time was active for 27 days and had 2 comments.

The density of the graph is 0.000119. The reciprocity ranges between 0.000459 and 0.0611 and the transitivity ranges between 0.000537 and 0.0123. The degree centralisation is 0.0562, the outdegree centralisation is also 0.0562 and the indegree centralisation is 0.000702

8.1.4 Subreddit Comparisons

It could be expected that the subreddits for DNM users and Clear Web marketplaces such as E-bay are used for similar purposes: to discuss product quality, the reputation of vendors, shipping times and costs, etc. however the fact that DNM subreddits include discussions of illegal activity may cause differences in the community structures of these subreddits and */r/Ebay*.

For example, we would expect a greater proportion of threads and usernames to be deleted or removed from the DNM focused subreddits. Further, we would expect more users of the DNM subreddits to use multiple accounts to prevent their activities being linked. This would lower the overall density, reciprocity and transitivity of the network.

The measurement values of each subreddit are presented in Table.8.1.

To evaluate if a greater proportion of content in the DNM subreddits is deleted, the Chi-Square test was used to compare the proportion of deleted comments and contributor accounts in the subreddit */r/Ebay* and each of the DNM subreddits. The proportion of deleted posts was not evaluated because of the high proportion of missing content information.

Comparing the proportion of deleted contributors in the subreddit */r/DarkNetMarkets* to the proportion of deleted contributor accounts in */r/Ebay* gives a test statistic of 0.0429 and p–value 0.836, therefore the hypothesis that the proportion of deleted accounts in the subreddit */r/Ebay* is distributed in the same manner as in the subreddit */r/DarkNetMarkets* should be rejected. The comparison of */r/Ebay* to the subreddit */r/dnmuk* gives a test statistic of 0.0191 and p-value of 0.890, therefore the hypothesis should again be rejected.

The Chi-Square test statistic on the comparison between the observed proportion of deleted comments in the subreddit */r/Ebay* and the expected proportion, calculated as the observed proportion of deleted comments in the subreddit */r/DarkNetMarkets*, was

TABLE 8.1: Summary of Subreddit Characteristics

Measure	*/r/DarkNetMarkets*	*/r/dnmuk*	*/r/Ebay*
Number of Posts	324,120	168,873	54,606
Number of Comments	572,585	281,248	87,264
Number of Contributors	417,701	200,792	64,428
Number of Posters	320,465	167,919	52,329
Number of Commenters	103,227	33,852	14,012
Deleted			
Posts	4,088	549	759
Comments	55,020	20,134	5,525
Contributor Accounts	69,858	27,643	5,798
Number of Comments Per Post (max, # 0's)	1.77 (895, 625)	1.67 (114, 204)	1.60 (25, 2,120)
Number of Posts Per Contributor (max, # 0's)	0.776 (139, 97,236)	0.841 (24, 32,873)	0.848 (1,606, 12,099)
Number of Comments Per Contributor (max, # 0's)	1.37 (20,040, 314,474)	1.40 (6,925, 166,940)	1.35 (3,657, 50,416)
Lifetime of Posts (max, # 0 days)	0.121 (30, 311,008)	0.0839 (28, 163,838)	0.0712 (27, 53,744)
Lifetime of Contributors (# 455 days, # 0 days)	5.22 (4, 400,017)	2.44 (3, 196,945)	7.33 (7, 60,798)
Density	$1.73 \cdot \exp(-5)$	$4.95 \cdot \exp(-5)$	0.000119
Reciprocity (range)	0.00122 - 0.0280	0.00141 - 0.0629	0.000459 - 0.0611
Transitivity	0.000329 - 0.00561	0.000368 - 0.0117	0.000537 - 0.0123
Centralisation	0.0515	0.0345	0.0562
Outdegree Centralisation	0.0477	0.0344	0.0562
Indegree Centralisation	0.00379	0.000689	0.000702

0.0124 with a p-value 0.911. These values for the comparison to the subreddit */r/dnmuk* were 0.00103 and 0.974, respectively.

Therefore, the proportion of deleted contributor accounts in the subreddit */r/Ebay* is lower than in either of the DNM subreddits, to a statistically significant level. However only */r/DarkNetMarkets* has a higher proportion of deleted comments.

The density of the subreddit */r/Ebay* is an order of magnitude higher than both DNM subreddits. This implies that, whilst none of the subreddits are highly connected, the subreddit */r/Ebay* was more connected. The Chi-Square test was once again employed to test the significance of this difference. It was used to measure if the ratio of edges drawn in the network constructed from the subreddit */r/Ebay* to those not drawn followed the same distribution as in either of the subreddits */r/DarkNetMarkets* or */r/dnmuk*.

The test statistics and p-values were 0.000598, 0.980 and $9.77 \cdot \exp(-5)$, 0.992 for each comparison and, therefore, the hypothesis that density of the subreddit */r/Ebay* followed the same distribution as either of the DNM subreddits is accepted.

This analysis supports the conclusion that users interact differently with Dark Web and Clear Web marketplace Reddit forums by leaving less of their comments and accounts undeleted but not that the amount of engagement is necessarily different.

8.2 Evaluation of Law Enforcement Interventions

The content and size of the Dark Web forums were analysed to understand how the forums were affected by the law enforcement interventions Operation Hyperion and Operation Bayonet.

For much of this analysis, the two Dark Web forums were combined and it will be explicitly stated where this is not the case.

8.2.1 Activity

An examination of the number of posts, comments and contributors over time shows that the activity over time is variable. The average number of posts per week is 4,780, the average number of comments is 8,420 and the average number of contributors is 6,200. Broadly, each variable follows the same pattern of fluctuations, implying that the community did not respond to events through a disproportionate increase of either posts, comments, or contributors but, rather, even in times of heightened activity the proportions of these variables relative to each other remained constant.

The amount of activity was observed over time and is given in figure 8.9. One such clear period of heightened activity is in the week succeeding Operation Bayonet on, approximately, the 13 July 2017. Here, the number of posts rises from 10,122 to 15,523, the number of comments rises from 17,829 to 27,803 and the number of contributors rises from 15,181 to 22,035.

To determine if this increase was disproportionate to the other fluctuations a moving average was employed to identify increases or decreases beyond 3 standard deviations from the mean. Where the window of the moving average was 15 or above, the number of posts, comments, and contributors on 13 July 2017 is more than 3 standard deviations from the mean, as shown in figure 8.10, 8.11, 8.12.

In contrast, across Operation Hyperion, the number of posts, comments and contributors falls (from 6,296 to 5,411, 11,282 to 9,651, and 8,683 to 7,882 respectively). At most, this value is 1 standard deviation away from the mean.

Whilst several other clear peaks can be seen just one (which took place in January 2017) reaches 3 standard deviations from the mean and only barely (in window 18). This is shown in figure 8.13, 8.14, 8.15. Therefore, across the measurement period, the largest increase in activity took place during the time of Operation Bayonet.

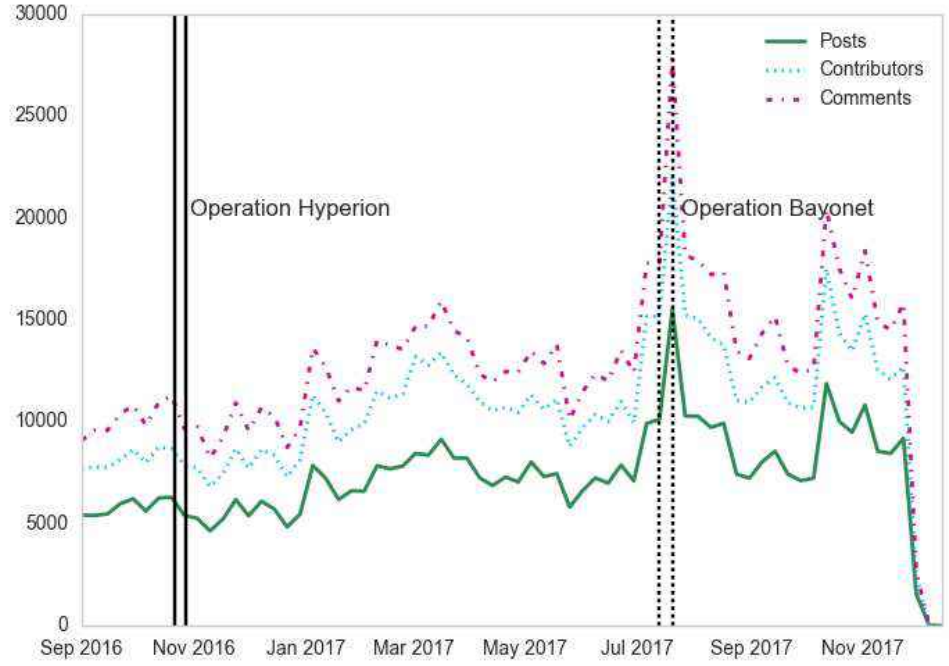

FIGURE 8.9: **Number of Contributors, Posts and Comments** in the DNM threads over the measurement period

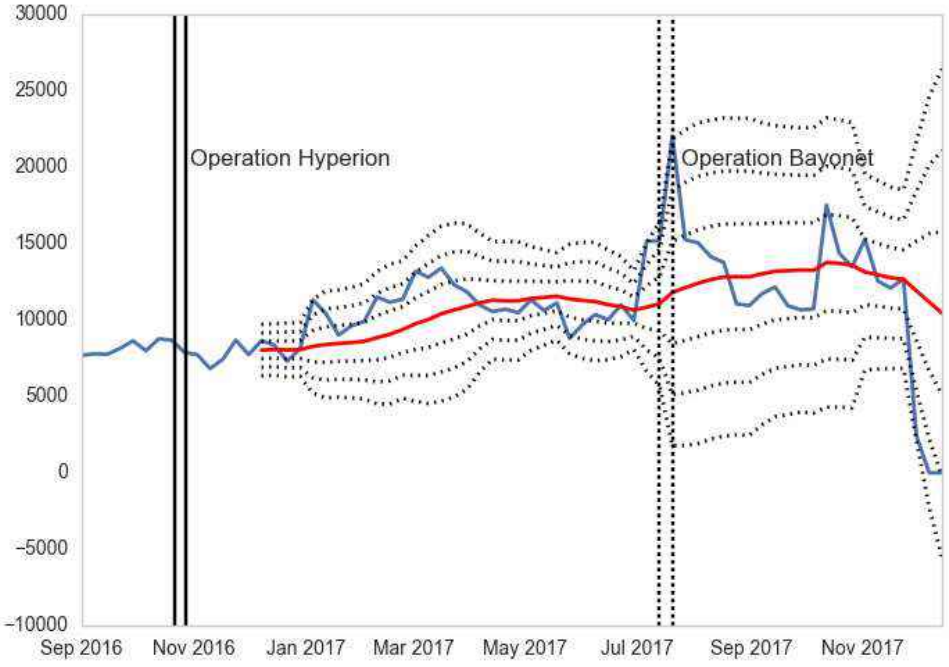

FIGURE 8.10: **Number of Contributors and Moving Average** with 15 Week Window.

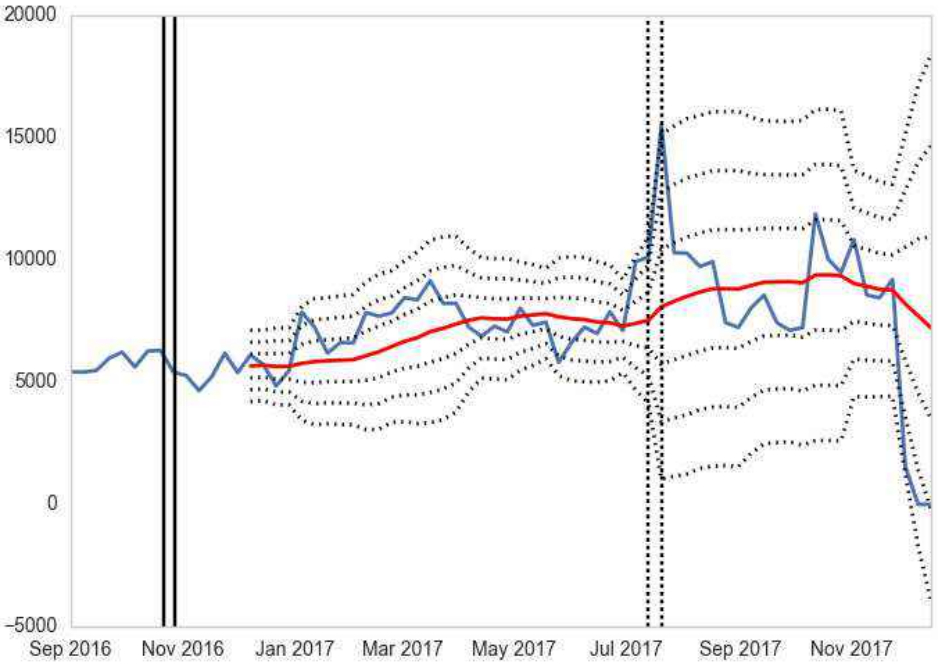

FIGURE 8.11: **Number of Posts and Moving Average** with 15 Week Window.

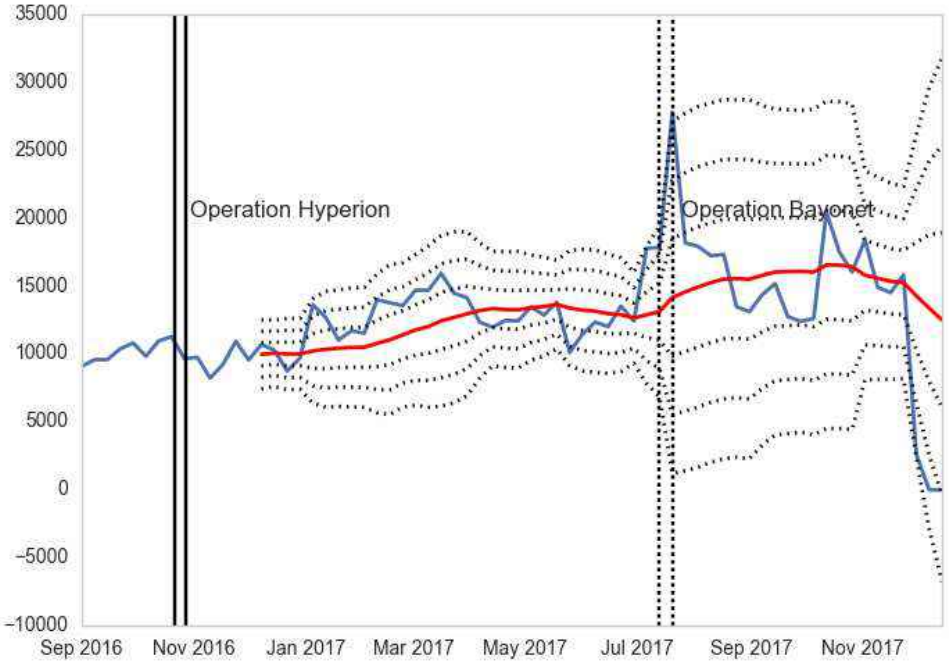

FIGURE 8.12: **Number of Comments and Moving Average** with 15 Week Window.

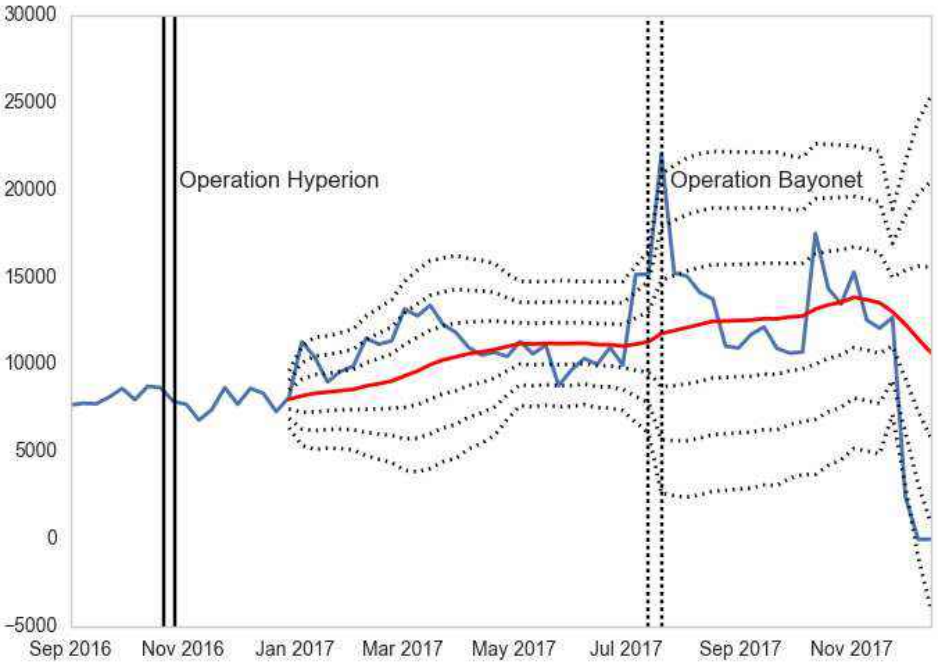

FIGURE 8.13: **Number of Contributors and Moving Average** with 18 Week Window.

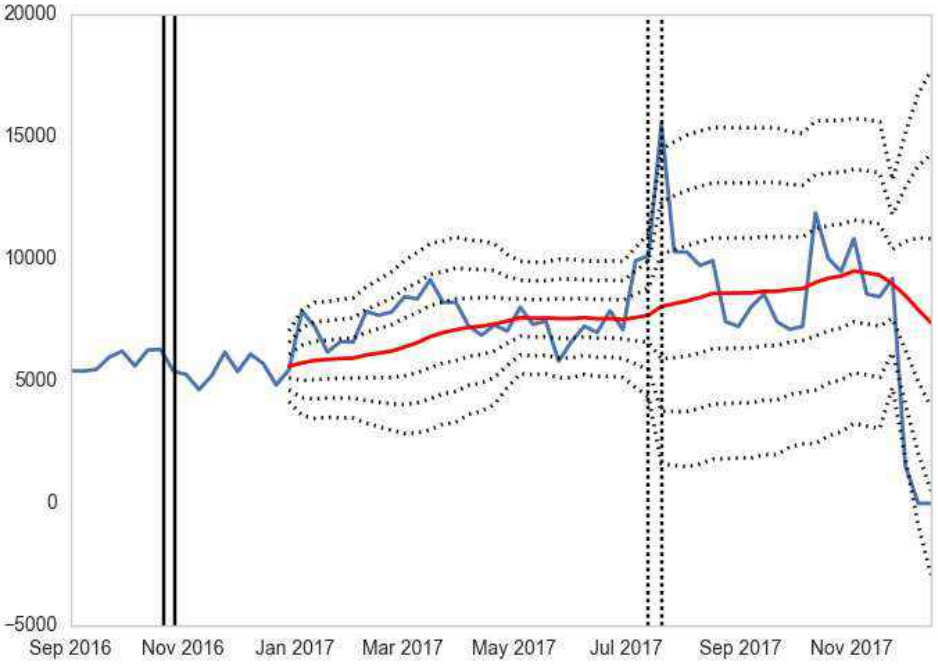

FIGURE 8.14: **Number of Posts and Moving Average** with 18 Week Window.

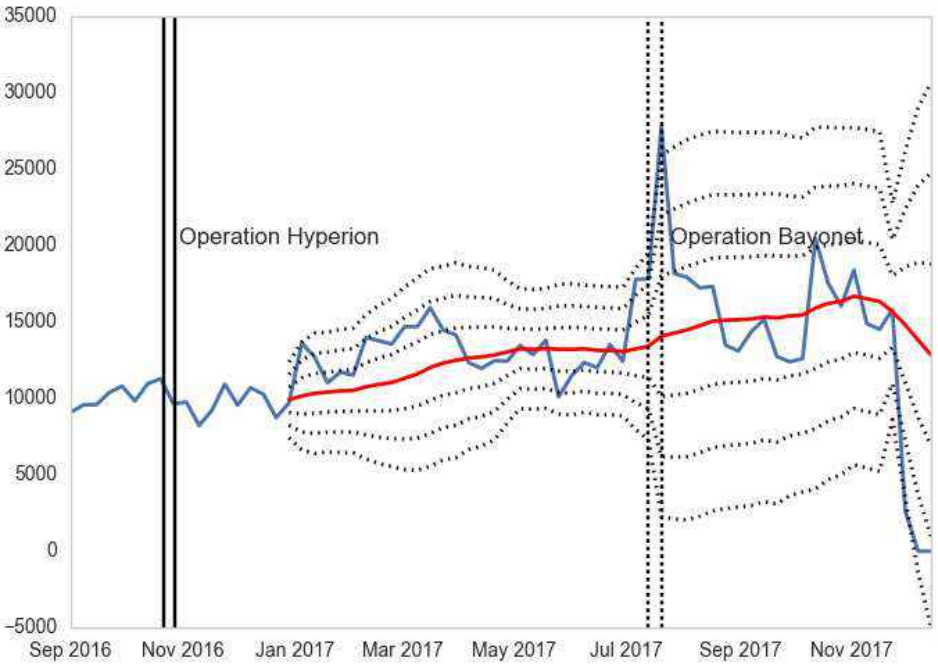

FIGURE 8.15: **Number of Comments and Moving Average** with 18 Week Window.

8.2.2 Content Analysis

The contents of the forums surrounding both Operation Hyperion and Operation Bayonet were explored qualitatively in order to understand how users felt about each operation, law enforcement in general, and the likely repercussions on themselves. Grounded Theory was applied to the posts and their comments, this approach involves an iterative process of reading the forums to identify key themes and/or concepts that are then substantiated through rereading and further data collection.

The operations were explored separately and their results compared at the end of the analysis to see if the same themes and perceptions emerged. As such, in the following sections, the processes of iterative understanding and data gathering will be described separately for each event.

8.2.2.1 Operation Hyperion

To begin with, relevant posts and their comments were collected from the week in which Operation Hyperion took place (22 – 28 October 2016). These were identified as posts that either contained a key word or posts with at least one comment that contained a key

‘arrest’, ‘love letter’, ‘ LL’, ‘a letter’, ‘police’, ‘nca’, ‘policeman’, ‘law enforcement’, ‘l.e.’, ‘ LE’, ‘fbi’, ‘operation’, ‘hyperion’, ‘confiscat’, ‘seiz’, ‘missing’

FIGURE 8.16: **Keywords used to identify relevant posts**

word. The key words were selected as terms used to refer to law enforcement explicitly (e.g. ‘police’ or ‘fbi’) or implicitly through law enforcement activity (e.g. ‘love letter’, a letter sent when a parcel has been seized) and are given in figure 8.16. These words were identified by reading through comments and posts before the initial analysis. Out of the 4,624 posts on the subreddit */r/DarkNetMarkets* and 729 on the subreddit */r/dnmuk*, this process identified 158 and 26 relevant posts for the subreddits respectively.

Of the 184 posts, 14 were deleted or removed from the subreddit and, of the 1,002 comments, 69 were removed or deleted. The posts covered a range of different topics including missing packages, vendor reviews, news stories about the Dark Web, OPSEC and how to respond to law enforcement warnings or package seizures. Of these, only 8 were considered to potentially be relevant to Operation Hyperion.

To be considered relevant, the posts or comments had to either refer to a letter explaining that the recipient was under suspicion (as opposed to a generic “love letter” or customs letter notifying the recipient of a confiscated parcel) or an increase in law enforcement activity. Generic discussions of law enforcement were disregarded if they weren’t considered to be talking about an actual event as were advice seeking posts in which buyers were trying to determine if their package had been seized but no evidence was provided that it had.

The first relevant post that was identified occurred on 24 October 2018 implying a delay in the reaction to the event. As such, the date range for data collection was extended by one month. In addition, the words “confiscate”, “seiz” and “missing” were removed from the set and “letter” was added to reflect the words that were found in relevant posts vs. irrelevant ones. For this new selection criteria, 896 posts were found from the subreddit */r/DarkNetMarkets* and 157 were found from the subreddit */r/dnmuk*.

Manual inspection of each discussion returned 78 relevant posts. The further from the date of the operation, the more likely relevant posts could only be identified through specific references to Operation Hyperion. As such, to screen the remainder of the dataset, all posts containing the term “hyperion” were collected. This returned an additional 38 posts, the last occurrence of which was on 20 July 2017.

Operation Hyperion was an international and multifaceted operation. Five different approaches employed by law enforcement were identified within the forum discussions parsed:

- Publishing lists of vendors/buyers who had been identified;
- Posting letters detailing suspected activity;
- Making phone calls or house calls to deliver similar messages;
- Seizing parcels;
- Making arrests.

These approaches were attributed to the law enforcement agencies from 5 different countries (Sweden, the Netherlands, Canada, the US, and the UK). They were discussed separately and then, later, simultaneously under the banner of Operation Hyperion.

The discussions covered a range of topics:

- posting news stories and discussing what has happened;
- speculating on how people were identified;
- identifying victims and speculating on the consequences for victims;
- giving advice on how to continue making purchases;
- sharing their opinions on the approach and on law enforcement.

.

Each of these topics will now be discussed in detail.

Description of the Operation.
When discussing Operation Hyperion as a whole, forum contributors posted links to news articles from a variety of different sources, including `ice.gov`, `stuff.co.nz`, `motherboard` and `cyberscoop.com`. Frequently, they would also post the contents of the article directly into the post so that readers would not have to click on the link ("*For those wanting avoid [sic] ice.gov...*"). These news stories described a multinational operation involving international partners and included official statements from a number of agencies. The agencies named were Europol, the NCA, the Australian Federal Police, the New Zealand Police and New Zealand Customs Service, Canada's Royal Canadian Mounted Police, Canada Post and Canada Border Service Agency, the Netherlands, French Customs National Intelligence and Investigations Directorate, Finnish Customs, Swedish Police Authority and Swedish Customs, Ireland's Garda National Drugs and Organised Crime Bureau, Spain's Guardia Civil, the FBI, the Five Eyes Law Enforcement Group and Interpol.

The news stories refer to multiple different tactics, including approaching suspects, sending letters and making phone calls. They also describe a website hosted by Dutch police which lists the usernames and suspected real world identities of buyers and vendors. This website is linked to in two posts (on 30 October and 29 November 2016) but neither post attributes the action to Operation Hyperion.

The first news story was posted on 1 November 2016. However, prior to this, posts and comments described the specific actions taken by law enforcement but did not explicitly link them to Operation Hyperion.

On 27 October 2016, the following was posted:

> i [sic] received a pretty unpleasant notice letter from NCA claiming that they have information that my address was used to send drugs via post from dark net before or on November 2014. It is just a notice, however if any more information is received a further action might be taken.

Similar letters were described both by other recipients and more generally ("*UK people getting letters*"). These letters were linked to Operation Hyperion on 7 November 2016 ("*I'm assuming our version was the love letters related to silk road*"). And letters were received as late as 29 November 2016 including letters that were not from the NCA and were not detailing purchases from 2014, as can be seen in the following exchange:

> I got a pesky LL from UK border force (pack was from CAN).

with the response

> operation hyperion ring any bells???

on 19 November.

This post alludes to a second approach attributed to the operation - an increase in seizures of Canadian packages.

> There have been many posts on packages not landing from Canada (including mine). It seems to have started during the Summer (Pangea?) but it seems to have carried on... Some have commented on receiving seizure notices from Canada on the market comments. Is it time for vendors to pause shipping to UK etc. I am sure vendors will be at risk to as they are being thoroughly investigated.

Operation Pangea is another international law enforcement operation specifically targeting online pharmaceuticals which had a week of action in June, 2016. Canada was a participating country and other comments which also refer to the low success rate of parcels shipped from Canada to the UK mention this operation.

Several posts also mention home visits from the FBI. Their descriptions of the interactions are similar to those of the letters received from the NCA, in that the purchases mentioned were from 2014. None of these encounters are directly attributed to Operation Hyperion by other contributors on the posts, however, several of the news articles posted state that the FBI approached people at their homes and, in a reply to a news story about Operation Hyperion, a contributor speculates "*Wonder if those posts that have been popping up regarding phone calls from the FBI have anything to do with this..*".

The amount of posts and comments in which contributors detail actual experience of Operation Hyperion are low, especially in comparison to the official figures posted (for example, an article claims the FBI approached 150 people (Cox (2016a)) and yet only a handful are found on the forum). This may be because suspects were unwilling to post about their experience on a public forum or felt they didn't need to having read other people's experiences. Given that individuals were targeted for purchases made in 2014, some may no longer be trading and so could have been out of touch with where to discuss the operation.

The different approaches and different law enforcement agencies only seem to have been linked after official statements were released implying that suspects targeted in the operation may not have known they were not the only ones without such statements. Instead, they may have felt that they were personally being targeted perhaps because of bad OPSEC, an informant or because they were of particular interest. The theories offered for how suspects were identified in Operation Hyperion are discussed in the next section.

Speculation on How Operation Hyperion was Conducted.
The DNM user experience of Operation Hyperion implied that law enforcement had uncovered evidence against their online activities. As a result, many contributors speculated as to how that evidence was uncovered. Several theories were posited which focused on the technical abilities of law enforcement, previous law enforcement operations, other DNM users giving evidence, using honeypot accounts or websites, and making up information.

As well as discussing how the operation worked, some contributors questioned the validity of the documents and phone calls people received as well as official documents published by law enforcement. When letters were first being received, contributors speculated that they were sent by vendors in place of products in order to scam their customers. Phone calls were thought of as potentially being from blackmailers. Elements such as the font ("*Fake, just the font gives it away*"), the header and footer ("*official letters always include a header and footer*"), and the way the law enforcement agency was described were all given as evidence ("*erm thers [sic] no such force as the*

uk counter narcotics team"). As more news articles about Operation Hyperion were shared, the theories about scams and fakes dissipated. All of the comments postulating the letters were fake were posted on 28 or 29 October and the last claim that the phone calls were fake took place on 6 November 2016.

As well as using the publications of Operation Hyperion to counter arguments attacking validity, the corresponding websites created by law enforcement were given as evidence, though these were also dismissed as fake by some ("*it's so hard to create a fake letter with a real URL on it isn't it...*"). Contributors also argued that the lack of a Bitcoin address meant that the letters weren't being used to exploit people.

> there would be no gain from sending a fake one because it isn't asking for anything usually fake letters ask for bitcoins

The dominant theory for how user details were obtained by law enforcement for Operation Hyperion was through the bust of Silk Road and/or Silk Road 2.0. This theory was often substantiated by the fact that the letters, house visits and phone calls referred to purchases made in 2014.

> I have a feeling that the address details were leaked when SR2 was busted, as these dates correspond, and I believe my address is not the only one leaked.

There is some speculation as to whether law enforcement were only able to obtain information from unencrypted messages though some letter recipients maintained that they did use PGP encryption on Silk Road 2.0 ("*I must admit I didnt use PGP right at the beginning but im very certain I did with SR2.*"). This was explained by law enforcement taking 3 years to decrypt the data (hence the delay in the operation) or somehow obtaining keys.

It also led to the theory that the addresses were obtained from arrested vendors who either gave up information on their buyers or stored buyer information in plaintext.

> You could have used PGP but vendor decrypted it, then saved it in to a word document, that's what happned [sic] in shiny flakes case

An early theory, before the extent to which people had been identified was known, was that a vendor was using buyer addresses as a return address. This information was presented as if it was a common method of identifying users ("*sounds like one of your vendors was using you as a return address*").

In discussions focused on the information advertised on the website hosted by the Dutch police, some contributors suggested that the information was collected through honeypot vendor accounts. This was evidenced by the small number of names which therefore could have feasibly been collected by one or two accounts.

> but that was a pretty tiny customer list on that site, honestly. I feel like a NL vendor (even w/o feedback) could easily get that many orders in a day or two, assumig [sic] they put up some good pictures and prices were reasonable. considering the nature of the darknet, they could even put full on fake-usernames on there and nobody would know the difference

The least popular hypothesis seems to be that information was collected using some kind of technical approach. One comment, which does not appear to have received any replies, claimed that information was found through analysis on the blockchain:

> I heard word that Operation Hyperion used this to bust people using new blockchain analysis... Basically LE defeated tumblers by watching movement of coins between exchanges and markets... LE did some follow up (you know, "routine investigation") and found that the person was a buyer or vendor.

Whilst this theory is not directly responded to, other comments discuss the perceived incompetency of law enforcement ("*the Dutch police force is probably the worst police force in the world, especially when it comes to cyber crime*"). This implies that contributors did not view Operation Hyperion as a demonstration of new law enforcement powers that make them vulnerable to arrest.

Instead, the discussions around how the operation take place preference de-anonymisation being the fault of those users who were identified. The users who received letters, phone calls or visits were isolated as those who did not use PGP encryption when sending messages on Silk Road or Silk Road 2.0. This behaviour would make sense as a method for mitigating the fear created by Operation Hyperion as users who believed they had good OPSEC may be confident they would not receive repercussions (and, if users were not identified in the operation, they make take this as confirmation that they have good OPSEC).

Potential Consequences for Suspects.
As well as speculating on how suspects were identified, contributors speculated on the potential consequences for those who received letters, or similar, or were identified on a website. Further, they discussed the consequences linked to Operation Hyperion as it was unfolding.

The key impact identified in Operation Hyperion, from the comments read, was an increased seizure rate of packages from Canada to the UK ("*So recently on DNMUK subreddit, there have been many posts on packages not landing from Canada*").

This led some contributors to claim they would stop ordering from Canada ("*Taken my faith from ordering from canada [sic]*") or warn others not to buy from Canada ("*UK is f***ed for weed.*").

Though this perceived increase was still questioned by other users

> Is there any actual increase of packages being siezed [sic]? Maybe its just one week in god knows how long that the postal service works with LE so they can show some packages to chiefs and have less to worry about for few weeks/months again.

And there were multiple claims of packages from Canada still making it to the UK:

> i [sic] have had stuff come from 4 of the active vendors ;) UK le know about a few of them and packages still come through.

Further, the impact of the operation was perceived to be short term. The statement

> I think I should be safe as Hyperion is over/relaxed now right?

was made on 3 March 2017.

Whether or not the supposed increased seizure rate was linked to Operation Hyperion, the fact that the two events coincided potentially allowed buyers who did not receive packages to recuperate their losses by proving their purchase had been intercepted using the customs seizure letters sent by law enforcement ("*Hopefully we get a letter bro at least we can prove it then*").

In terms of future consequences, the vast majority of speculators concluded that there would be none, however a few did claim that Operation Hyperion could lead to more serious, unspecified consequences for users

> they don't seem like they're playing around right now and they have been going after bulk buyers

and the DNM ecosystem as a whole

honestly the DNMS are gonna be dead soon. USA doesnt f*** around

When discussing what would happen to the recipients of the letters, contributors argued that the letters themselves demonstrated nothing more would happen ("*If they had more s*** on us they would have done something by now bin it.*") and substantiated this claim using the text within the letter itself ("*It confirms they are not pursuing criminal action so I wouldn't worry too much*"). Further reasoning came from the size of the operation ("*Think about how many users and orders there were... it'd take a while for them to talk to everyone.*) and personal experience of similar actions ("*It looks more of an advisory/deterrent letter. I got a similar one a few years ago*").

Despite such a high proportion of comments claiming that users who received letters or similar would face no repercussions, some comments still asked for or offered advice on how they could best protect themselves. These are discussed next.

Advice Given to Suspects.
Precautionary measures were recommended for users who had been targeted by Operation Hyperion. These focused on changing address ("*It would be a good idea to move or get a new drop before ordering again.*") as this was potentially how they were being monitored ("*But i would now be very very carefull if ordering anything else just incase they monitor your mail*"). However, these suggestions were treated with derision by other contributors (*Yeah maybe a scare tactic but I'd move if it were upto me LOL brb moving house because i received a letter.*

Several contributors made hyperbolic suggestions, potentially to mock law enforcement or other users who considered Operation Hyperion to pose a threat ("*Send them back a letter denying everything. But, PGP encrypt it*", "*Clean house, smother the tails usb in cement and swallow it*"). A trope used to mock contributors to the forum concerned about other law enforcement interventions (such as customs seizure notices) in which the contributor is advised to move to Belize was also employed (*your ass is on a one way flight to Belize*).

This implies that contributors to the forum either felt that it was not necessary to take precautionary action despite the operation or that it was unclear what could be done. The former is more likely as, it was pointed out by a contributor, some of those targeted by the operation continued to receive packages to their address after 2014 (the time period during which the letters and phone calls listed their activity) and so it was unlikely law enforcement were able to act on their information.

> youve [sic] also been having packages sent from nov14-july16 without any problems. order some weed smoke a j and chill out ;) no judge would give a warrant based on an address 2 years ago as long as the stealth is adequate

It is clear that many contributors claimed to be unconvinced by the potential for serious repercussions from Operation Hyperion. The wider variety of opinions on the operation, and the law enforcement teams who participated in it, will now be presented.

Opinions on Operation Hyperion.
A large portion of the discussion of Operation Hyperion was in the form of contributors sharing their thoughts and opinions on why it had happened, its implications on the future of the DNM ecosystem, and the ethics of the approach.

Several contributors felt that the operation was a waste of resources or an illegitimate action. Some argued that there were more important issues, such as terrorism and child pornography, to tackle, and others that drugs should be legalised because of the demand for them and the rights of individuals to put them in their bodies.

> Drug dealers/vendors are not the problem. If the police arrest every dealer, new ones will sprout up like weeds one day later because of the extremely high demand of every type of drug out there. Supply, demand... Legalize

Following from this rhetoric, some argued that this operation was unacceptable ("*this isn't protection. This is tyranny. This is a violation of human rights*").

Unsurprisingly, no one was sharing opinions that the operation was a good thing that would benefit society, though some comments alluded to this sentiment being shared on other social media platforms. Instead, contributors to this forum argued that Operation Hyperion, rather than being a display of law enforcement power, was a tactic designed to scare users because there were not enough resources to take more effective action.

> Seems like theyre [sic] scraping the bottom of the SilkRoad 2.0 barrel and making a big noise in the hope that they will keep their funding

> You'd imagine a world-wide super counter-narcotrafficking union to the Z power would have much more names

> That's 20,000 people they have to monitor, now. How far will their resources stretch? How many people are they willing to go after, for small amounts of whatever? LOL. The joke's on YOU, law enforcement.

This was particularly targeted at the FBI who was seen to have a lower impact than some European organisations.

> I don't understand this move by the FBI, other than being jealous of the big busts in europe [sic]. If their goal was to "send a message", the message sent was, "If the FBI find out you're buying drugs online, you'll get a stern warning."

The intentions of law enforcement were further undermined by speculation on why the operation had been so publicly advertised. Some claimed that it was "counter productive" and that it was beneficial to DNM users ("*Also gives vendors a heads up they are being watched. Which means time to rebrand and reevaluate*"). Others added that these kinds of operations were responsible for the growth of the ecosystem through advertising it to the world ("*It's the f***ing FBI who made the dark markets explode by making a high profile case of the SR one bust*").

Though, three comments did articulate some intimidation:

> Quite an intimidating site IMO. If I seen my market username there I'd s*** myself

> Although I must admit I am happy I am not on their bloody list

> this Operation Hyperion s*** is hitting hard.. I am seeing a ton of busts

Several comments implied that contributors expected some form of large scale law enforcement operation, annually. They compared Operation Hyperion to previous approaches and found it preferable to the shut down of large or multiple DNMs.

> Is this all worldwide LE has for DNM busts this year? If so they ended this year damn near empty handed: They're grasping for anything, going after buyers with sloppy opsec. Looks like we're winning, slowly but surely. Let freedom ring!!!

> Operation Hyperion. It's that time of the year folks. Almost every year at this time, since the fall of our beloved SR, we see a spike in DNM leo activities.

If the comments made by contributors on the topic of Operation Hyperion truly reflect their opinions and are indicators of how they behaved, Operation Hyperion did make some users concerned for their safety and the stability of the ecosystem. However, the conversation implies that the majority of contributors felt no tangible repercussions and interpreted the operation as a sign of weakness on the part of law enforcement, particularly in countries where no arrests were made. Further, where consequences were felt (for example in the perceived impact on trade from Canada), these consequences were considered temporary and treated as an expected event within the ecosystem calendar.

Summary

The posts and comments that could be attributed to Operation Hyperion discussed the nature of the operation, its repercussions and its validity. Contributors shared several different theories about how the operation was conducted and appear to preference the theories that claim users were identified because of errors in their own OPSEC, rather than the technical abilities of law enforcement.

Whilst several contributors articulated the view that law enforcement operations against DNMs are illegitimate, many felt that these operations were inevitable. When compared to other operations, Operation Hyperion was not considered to have a big impact. Though some contributors shared advice on how letter recipients could protect themselves from law enforcement, many felt there would be no actual repercussions.

8.2.2.2 Operation Bayonet

To find posts and comments relevant to Operation Bayonet, the terms "alphabay", "ab" (a common abbreviation for *Alphabay*), "hansa", and "bayonet" were searched in posts from July 2017. Out of a total of 46,130 posts, 2,848 contained relevant key words. These were each read to find those that were relevant to Operation Bayonet.

The posts prior to official statements about what had happened to *Alphabay* and posts after this date were analysed separately to see if users responded to the event differently when they knew for certain that law enforcement were involved. Official statements were released on 13 July and it was also after this date that the site *Hansa* went offline. Therefore, how users responded to this second site disappearing was also examined.

Posts and Comments Between 1 July and 13 July 2017

Alphabay was first reported as down on the 5 July 2017 and this news received varied responses. Some users immediately commented that *Alphabay* was down permanently ("*dnstats shows it's down as is the forum, is it over for AB?*") whereas others cited *Alphabay's* history as evidence that the shutdown being only temporary ("*You aren't used to this yet? Lol AB has exit scammed what 67 times in the past six months.*").

News of the alleged owner of *Alphabay*, Alexandre Cazes's, arrest and death was not posted until the 13 July. In the 8 days before official information came to light, contributors speculated on the current state of *Alphabay*, but also continued to participate in other types of conversations on the forum. Even in the posts which specifically mention *Alphabay* or *Hansa*, topics not directly about the *Alphabay* closure arose.

For example, contributors continued to discuss vendors, their products and shipping times ("*Might not be the best time for this, i'm [sic] sure others are worried about things right now lol but im [sic] my experience anywhere between 2DD-14DD - so chill out and wait i [sic] guess*") though these discussions were now often focused to how the site closure would affect their purchase; *Alphabay's* tumbling service ("*I don't think they shouldn't take fees dude. They take 3% out of every sale... This is a criticism of their tumbler which takes way more money than is necessary.*"); and, other marketplaces in general and as a replacement to *Alphabay*. Additionally, the death of a student linked to *Alphabay* was mentioned ("*Seems most likely LE seized her computer which was unencrypted so easy to access... Since the order was made under a week ago it'd still be in AB's order history.*").

Posts about *Hansa* debated the multisig element of the site, whether it would protect against exit scams and how feasible it was to use.

These conversations consisted of 43% of those that were read demonstrating that *Alphabay* going down was not the only topic of discussion in the forum.

The closure of *Alphabay* was discussed in a number of ways. Contributors posed theories on what had caused the site to go offline; debated whether the downtime was permanent or temporary; shared how much they had lost after the site went down; discussed whether or not to continue trading and, if so, which DNM to switch to; and speculated on the potential consequences for *Alphabay* users.

Theories

Approximately one third (413) of all the coded comments were part of discussions about what happened to *Alphabay*. The theories posed included *Alphabay* going temporarily offline for site maintenance, for example to implement Zcash or to fix a technical issue (13%); being hacked, doxxed, a victim of malware such as ransomware or otherwise attacked (2%); being shut down by law enforcement (either Russian, Canadian, or German) who had found a Tor exploit or arrested a site admin in a raid (38%); closing in an exit scam (46%) or an exit run (1%) if law enforcement had seized the site servers; or being closed by the Russian Mafia (<1%).

Examining the number of mentions of each theory over time shows that the hypothesis that the site was down for site maintenance decreased over time. Whereas the view that

it was an exit scam or law enforcement intervention either remained consistently high or increased over time.

The theories were substantiated with, or potentially grew from, news stories and observations of other relevant websites and tools, for example the Blockchain. On 5 July, someone noticed that a large amount of BTC had been transferred from what was believed to be wallets connected to *Alphabay*. Its worth was estimated at being between $3 and $4 million. Some argued that this was evidence of an exit scam ("*It's strongly looking like AB has done an exit scam for 4 million $*") however others felt that it was too small an amount and, instead, the money had been transferred, for example, to help with the inclusion of ZCash ("*3 million isnt [sic] a savings pot to them, its an investment fund. most likely for zcash*") it was then later (on 7 July) argued that the Bitcoin wallet was not associated with *Alphabay* at all ("*it wasnt [sic] the wallet of alphabay, its [sic] already proven. A bitcoin selling page got hacked, thats [sic] where the 4m$ come from.*").

Contributors hypothesised that recent raids in Canada were potentially linked to the *Alphabay* down time, however, it was widely believed at the time that *Alphabay* was hosted in Russia and that its admin were Russian, therefore the theory was dismissed by many ("*Guys, Alphabay is very obviously based in Russia, not the eastern part of Canada... Alphabay also has multiple servers and I feel they would be to [sic] smart to host them in a country like Canada when there are dozens of better countries with better privacy laws*"). A similar process was used to link the closure to a German law enforcement operation ("*there was a big pedophile [sic] ring busted in germany [sic] today. maybe the'yre [sic] related?*").

As well as posing theories, contributors analysed and critiqued the theories within the forum. There were 62 comments demonstrating scepticism of the theories for why *Alphabay* was offline. Contributors argued that an exit scam was unlikely because of the amount of money likely stolen by the admin team ("*Was worried too at first but do the math in the amount.. that's like what, less than 350k? If the biggest DNM exit scammed, surely it would be for millions?*") and how it didn't fit the pattern of other exit scams, e.g. when Evolution disabled withdrawals prior to shutting down ("*From an experienced vendors perspective - We have seen many exit scams and there is a common theme, it always starts with withdrawal problems.*"). The exit scam theory also became less popular as contributors became more convinced of other theories ("*Why do you think they exit scammed when all the evidence leads to them getting busted?*"). Figure 8.17 shows how the number of comments that alluded to each theory changed over the course of the observation period.

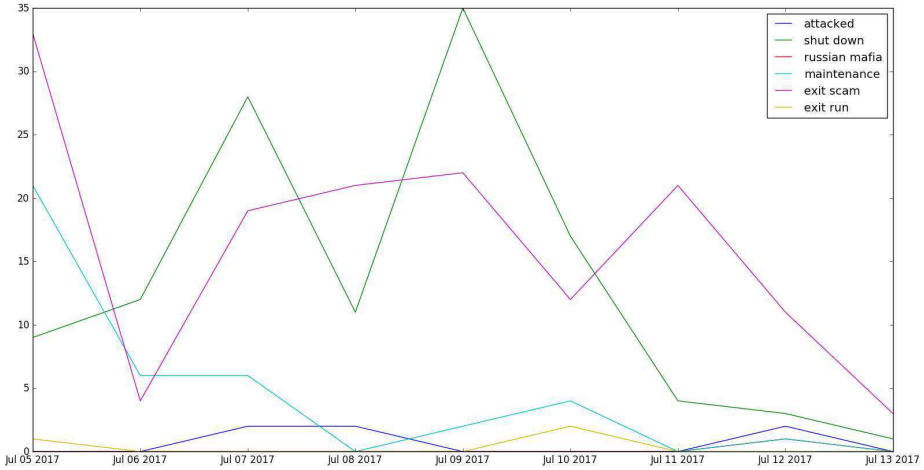

FIGURE 8.17: **Number of Mentions of Each Theory** Over Time.

Contributors were mostly sceptical of law enforcement involvement because of the lack of official word from any agency ("*If FBI seized it wouldn't they upload their banner thing saying it's been seized?*", "*I assume LE had control over alphabay servers they would be bragging about it.*"). Another criticism of the theory was that *Alphabay* could not so easily be taken down by law enforcement ("*AlphaBay being "seized" seems unlikely to me since it's fairly clear that AB is not hosted on just one server. They've got 8 onion addresses... I'm no computer wizard, but I just can't believe their setup is that bad.*").

However, the lack of communication from *Alphabay* moderators was given as a reason why the site was not simply down for maintenance ("*Something bad happened. Telling you guys they are updating AB without a prior warning is bulls***.*"). In addition, when contributors posted on the subreddits claiming to be *Alphabay* admin or moderators, their credentials were heavily scrutinised and their comments were entirely disregarded by some ("*'i [sic] dont [sic] trust this*").

On the 8 July some vendor stats appeared on people's *Dream* accounts and this was seen as a positive sign of *Alphabay's* return ("*It is reassuring that they could access our Vendor stats on AB, they do show me AB feedback. Hopefully that means part of AB is still running somewhere or somehow. Not sure*").

In these discussions it was also speculated as to whether the downtime was permanent or temporary, these opinions are described next.

Permanent vs Temporary

417 comments explicitly indicated if the contributor considered the site to be down permanently or temporarily. 93 (22%) of these indicated that the contributor considered

the site to be down permanently, either by explicitly stating the site was gone forever ("*Alphabay won't return. Period. Don't have false hope.*"), by actively seeking for a replacement to *Alphabay* ("*What are some markets we can migrate to?*") or referencing the way they felt *Alphabay* had closed ("*So why not cash the f*** out when quitting instead of doing the "right" thing?*"). A further 27 (6%) comments did not explicitly conclude that *Alphabay* was permanently shut down but did imply they felt it was more likely, i.e. were pessimistic about *Alphabay's* return. These comments expressed fear at *Alphabay's* fate ("*It's the lack of communication that has me worried.*") or otherwise indicated the contributor was not confident of its return ("*Not looking too good boys. . .*"). An additional 46 (11%) comments referred to *Alphabay* in past tense and so, even though they did not state that *Alphabay* was offline permanently, were considered to show that the contributor did not believe *Alphabay* would come back.

76 (18%) comments explicitly referred to *Alphabay* coming back online for example by referencing the reason they thought the site was currently offline ("*they are implementing Zcash, nothing to worry about fools*"), dismissing the behaviour as expected ("*You aren't used to this yet? Lol AB has exit scammed what 67 times in the past six months.*"), or by talking about what would happen when *Alphabay* returned ("*Once Alphabay comes back up I'll have a gander then pal.*"). A further 54 (13%) comments referred to Alphabay in the present tense e.g. by advising other contributors to find vendors on *Alphabay* ("*You can find them all on TMG, Dream, AB and DHL*") and/or by referencing *Alphabay* in the same way that active DNMs were referenced. These contributors were also considered to view the downtime as temporary as, otherwise, it would be expected that they use the past tense to refer to the site.

As well as this, 24 (6%) of the comments centred around asking users not to spread "FUD" (fear, uncertainty and doubt). This is a concept used within the forums to denounce fearmongering and speculation and highlights an awareness of how users may be deterred from participating on the ecosystem even when they are not at risk. Whilst these comments did not overtly argue that *Alphabay* would return, they criticised comments arguing it would never return for being hyperbolic ("*Paranoia from drugs most likely mate if u ain't selling chill*") or unsubstantiated by evidence. It was therefore assumed that the contributors felt that *Alphabay* would return.

Two conditional emotions were expressed towards *Alphabay.* The first was of contributors hoping that *Alphabay* would return, this was expressed in 52 (12%) of the comments. In these comments, contributors stated that they hoped *Alphabay* would return, most commonly because they had coin stored on the site or were awaiting an order ("*Made my first order with meerkovo yesterday, don't even know if order even got accepted wish it would come back up*"). These contributors did not state that they thought *Alphabay*

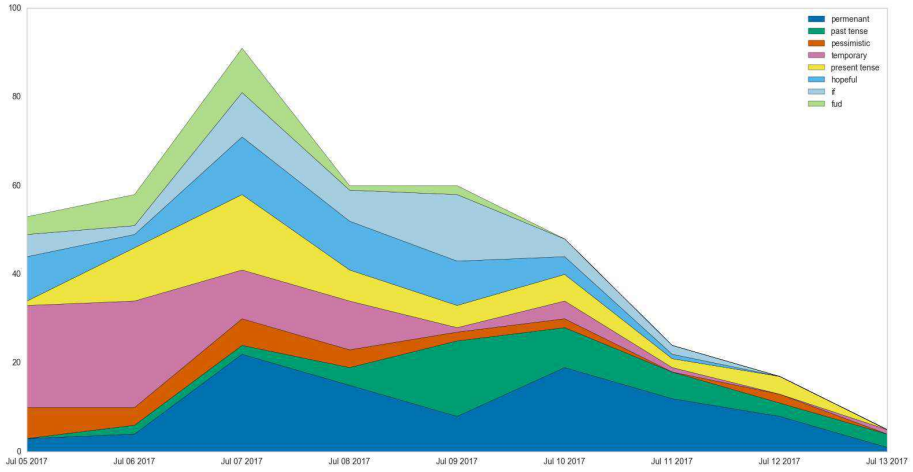

FIGURE 8.18: **Number of Comments** Containing Each Perspective on if Alphabay was Down Permanently or Temporarily

would be offline permanently or only temporarily. Similarly 45 (11%) of the comments referred to what would happen if *Alphabay* went offline permanently or if it came back up, without commenting on what they thought would actually happen.

Figure 8.18 shows the number of comments for each perspective during the observation period. This shows that perspectives which considered the closure to be temporary were more common at the start of the downtime and diminished over time, as did comments that referred to the downtime being permanent as FUD.

Many of the comments that considered *Alphabay* to be closed were in the context of users discussing what they had lost as a consequence. These impacts of *Alphabay's* closure are now presented.

Losses

There were 164 comments in which contributors shared what they lost from *Alphabay* shutting down and an additional 31 comments in which contributors sought out vendors with whom they had traded with on *Alphabay* in order to conclude their purchases.

These comments identified the following losses that the users of *Alphabay* experienced:

- 110 (67%) comments described losing money held in escrow ("*I have over £20k in escrow on AB. I know of another UK vendor with over £50k in escrow on AB :(*");
- 37 (23%) comments described losing at least one product ("*had an alphabay order that was shipped priority on the* 29^{th} *and it's not here. Of course my vendor is probably too stressed out herself to find out where it is.*");

- 5 (3%) comments described losing reputation or feedback ("*I didn't lose to [sic] much money ($700). What I lose [sic] was my great feedback on AB. It will be a slow couple of weeks building my prestige back up.*");
- 2 (1%) comments described losing the ability to contact a vendor or customers ("*We lost a fair bit of coin, but money can be replaced. I'm more bothered about how much custom will be received through Dream and Hansa*");
- 13 (8%) comments described losing nothing in the site closure ("*we lost absolutely nothing in the ab exit.*");
- 8 (5%) comments described losing something but did not specify what was lost ("*i've [sic] already been f***ed by alphabay this week, maybe next week*").

In 44 instances, whether the contributor felt that they had lost a large or small amount could be estimated. 25 (57%) such contributors claimed to have lost, or be in the position to lose, a substantial amount ("*I have enough money to buy a decent car on the line.*", "*anyway lost a f***ing fortune thanks to this*", "*I have over 320k in escrow on AB.*"). These discussions led some contributors to speculate if the *Alphabay* closure would be enough to put users out of business:

> Yeah, it'll sting a little but I'll get over it. I'm well stocked with everything so it's back to business as usual for me. I know of another UK vendor with 50k in escrow and he'll be just fine too. Some small vendors will not be so lucky. If a small guy just reupped and sold most of your stock in the week before AB went down they'll be in a tough spot having lost all their product and their coins :(

Contributors who lost money were blamed for leaving money on their account ("*So this seems like a decent thread to ask in, can OP or anyone else explain to me why you leave considerable amounts of cash on the market?*") especially by users who claimed to have not lost anything or very little ("*I lost nothing because I withdrew my change*"). There was also an emerging theme that exit scams were an inevitable element of trading on the Dark Web and so the costs should be managed accordingly ("*you know you're getting BTC taken at some point of [sic] you are in escrow. S*** happens and you move on. Can't complain really*").

Additional harms of the closure included users impersonating established vendors in order to scam their customers ("*Regardless of the vendor, remember to always make the vendor verify themselves by using a known key of the vendor. Unfortunately during these times, scammers come out of the woodwork and will impersonate well known vendors. We had to take a few down last night (not for this vendor).*", "*Been looking at LC on Hansa but still leery about anybody right now,,[sic] afraid I'll pick someone trying to*

*make up for lost AB coin or some s*** at my expense.*") and phishing links described as the *Alphabay* website. 14 comments discussed more than one fake *Alphabay* site created in order to steal users' log in credentials and money they were prepared to transfer to the site.

Contributors who believed that *Alphabay* may have been seized by law enforcement were also concerned for their security. The consequences of law enforcement seizing an *Alphabay* server were predicted to be anything on a scale of no consequences ("*If simply one server was seized and Admins have good opsec why do they need to quit while they're ahead?*") to extreme consequences ("*The dark nets are getting scary these days - f*** I hope this isn't the beginning of the end !?*"). Contributors also offered advice on how others could protect themselves from possible law enforcement investigations. Similary, these recommendations ranged from the very extreme ("*Word of advice to everyone that thinks some s*** might go down . Go get you a good lawyer on retainer pay all upfront , clean house , have a nice bond money set aside . This s*** is very real*") to the very small ("*Password change is a must, yes.*"). In addition, some users felt that only those users who had relied on *Alphabay's* autoencrypt system or who had made large purchases were at risk of law enforcement action ("*Sending coin from coinbase to AB was a mistake. Using the market to encrypt your address was a mistake. Maybe you didn't run into a problem doing it that way yet, but you are the low hanging fruit out of all of us here if you keep doing things that way.*", "*Yeah, there's thousands of buyers on AB I doubt they'd go after any unless you bought like a kilo of coke or pills.*").

To Trade or Not to Trade

Another discussion topic was whether or not it was safe to continue trading on the ecosystem. 132 comments indicated that the contributor would continue to trade, just 2 indicated they would cease trading and 7 indicated that the contributor intended to take a break before continuing to trade. There were also 94 comments in which the contributor queried which DNM would provide a suitable replacement to Alphabay as well as 15 comments explaining how the contributor would continue to trade but not on Alphabay, even if it returned.

The reasons given for not continuing to trade were the contributors' bad luck with exit scams and not being able to find the right product for the right price on remaining DNMs ("*I do, but not because I lost money. I only order domestic and all my favourite vendors were on AB, not a single one on Hansa or Dream. Only offer on Hansa for weed is 70$ for 5g, no thanks.*"). The reasons given for taking a break before restarting on another DNM were a lack of funds ("*i've already been f***ed by alphabay this week, maybe next week*"), the ecosystem being in turmoil ("*Before AB went down I did have larger scale testing plans with Chemicals_Spain and Meerkovo but I'll wait for calm before starting*

that again.") and taking precautions in case law enforcement were involved in the shut down.

The market most recommended or mentioned as one to move to was *Hansa* which featured in 59 comments. 3 contributors said they would not move to *Hansa*, one because it had a bad user interface and two because it was preventing new sign ups. Whilst most contributors did not give a reason for why they had chosen *Hansa*, the most cited reason was that *Hansa* was safer because its multisig feature would protect users against exit scams.

In contrast, the next recommended DNM, *Dream* was warned against in nearly as many comments as it was recommended because contributors were worried it was likely to exit scam in a similar manner to *Alphabay*. Contributors who did favour *Dream* cited its wide variety of products and vendors which was considered to be greater than that of *Hansa*.

The other DNMs that were discussed were *DHL*, *Traderoute*, *Wallstreet*, *Zion*, *CGMC*, *Valhalla*, and *RSClub*. One vendor also mentioned a listing they had on E-Bay and one buyer declared that they would now prefer to conduct direct deals with vendors, over using DNMs ("*Probably just gonna direct deal with most of my vendors from now on, can get a better price anyway.*").

Contributors who stated they would not use *Alphabay* if it returned either said this was because of the poor customer service during the downtime ("*If Ab does come back, Im [sic] spending what little I had there and thats that. AB have been extremely unprofessional, not communique, no official posts, nothing.*") or because they were concerned of law enforcement involvement ("*i [sic] doubt i'll [sic] us [sic] AB again if it comes back. honey-pot seems too likely.*").

Summary

Before contributors knew for certain of law enforcement involvement, the majority did not believe *Alphabay* to be gone permanently. Of those who did, many were content to move to another DNM to continue trading and considered the consequences of the event to be small. Though some contributors suffered losses, other members of the community were quick to judge them for keeping their Bitcoin on *Alphabay*. The lack of information did lead contributors to speculate about law enforcement involvement and imagine grave consequences however, often, competing theories with less severe consequences were favoured.

13 July 2017 Onwards

There were 2,140 posts and 8,421 comments between 13 July and 31 July 2017. 804 of the comments had been removed or deleted and a further 2,113 comments could not be labelled because they did not contain enough information or context to determine their contents. 327 posts (containing 1,197 comments) and a further 47 comments were labelled as not relevant. These posts and comments discussed 13 different topics:

- Advertisements for products (13 posts, 42 comments) e.g. "*Come find us on TradeRoute! All quantities available (qp, hp, full p).*"
- Bitcoin related posts (13 posts, 67 comments) such as the fork scheduled for August, 2017 and how it might affect the ecosystem ("*I keep seeing s*** about BTC forking. Is there an actual chance markets may be disrupted for a while ? That is scary because this is the only way for me to make actual money.*")
- Busts and arrests not considered related to the operation (7 posts, 122 comments) such as the arrest of a vendor found to have "*. . . several hundred grams of fent? Probably coming from China. Small packages from China must be getting extra special love at the border. The seller could be flagged as well.*"
- Buyers and vendors airing disputes from a marketplace or discussing the dispute system (2 posts, 6 comments) ("*Disputed with Turpin, won and got my money back.*")
- Discussions of drug usage/habits (24 posts, 184 comments) for example users talking about their favourite products, managing addictions or plans to use drugs in the future ("*i recently bought 20 tabs of lsd just for personal use, but yeah i guess it looks like a small time dealer to police*").
- A conversation about how the ecosystem search engine Grams operates (1 post, 2 comments).
- General queries about markets in the ecosystem (42 posts, 148 comments) for example, which one to use, how they compare, and what services each provides ("*I'm not sure it is 100% multisig. Zion is. Wall Street is 100% no account deposits. Like hansa was. Well I don't think anyone is really even using wall street. Dream has 100000 listings. TradeRoute has 10000 listings. Wall Street has 1000 listings.*")
- Noob questions (1 post, 17 comments) the subreddit */r/DarkNetMarkets* allowed one post a week in which users could ask basic questions about the ecosystem, for example, about how to create an account or perform basic opsec tasks.
- Queries on an order that had been placed (2 posts, 17 comments) for example asking when products are likely to arrive or complaining about the postal service and the time it takes for packages to arrive ("*Had five packs land this week, although still have quite a few outstanding*").

- Reviews or advice sought on a particular product (41 posts, 150 comments) for example contributors leaving comments about products they had tried are asking for recommendations of vendors from whom they could purchase particular products ("*Is it possible to find aripiprazole (Abilify) on the darknet and....where do I start? I'll take all the help I can get*").
- S***posting, i.e. comments and posts that are jokes or otherwise intentionally misleading or unhelpful, these were identified through other users calling them out as s***posting or similar (2 posts, 25 comments).
- Technical queries that were unrelated to improving opsec in the wake of Operation Bayonet (13 posts, 49 comments), these included general questions about VPN and multisig tools.
- Reviews or advice sought on a particular vendor (262 posts, 937), these posts and comments included leaving positive reviews about a vendor's product, opsec and/or customer service as well as discussing vendors more generally. For example, one post contained a discussion about a vendor whose customer doxxed themselves in a review.

9 themes emerged within the remaining posts and comments. These were as follows:

- **Alphabay admin** (120 posts, 240 comments): posts and comments about the admin of *Alphabay* and what has happened to them, mostly focusing on Alexandre Cazes who allegedly ran *Alphabay* under the usernames of DeSnake and Alpha02.
- **The consequences of the operation** (601 posts, 979 comments): posts and comments that either speculated or recorded the consequences of the operation as well as posts and comments about who the operation was likely to affect or not affect.
- **Whether or not users would continue trading** (148 posts, 212 comments): posts and comments in which contributors discuss if they will continue trading, the reasons they give and the markets they intend to use. These posts and comments each explicitly state or imply that the contributor was active on *Alphabay* and/or *Hansa* and therefore are not general comments about where to trade (which were categorised as 'other').
- **Hansa** (203 posts, 342 comments): posts and comments about *Hansa* when it was still active. This includes discussions of the market, the ban of fentanyl, and issues that people had with their accounts immediately before it went down.
- **Opinions** (211 posts, 382 comments): posts and comments in which contributors expressed an opinion about the nature of the operation, for instance whether or not they thought that it would be successful or if it was a legitimate action. This category also includes opinion expressed about law enforcement in general.

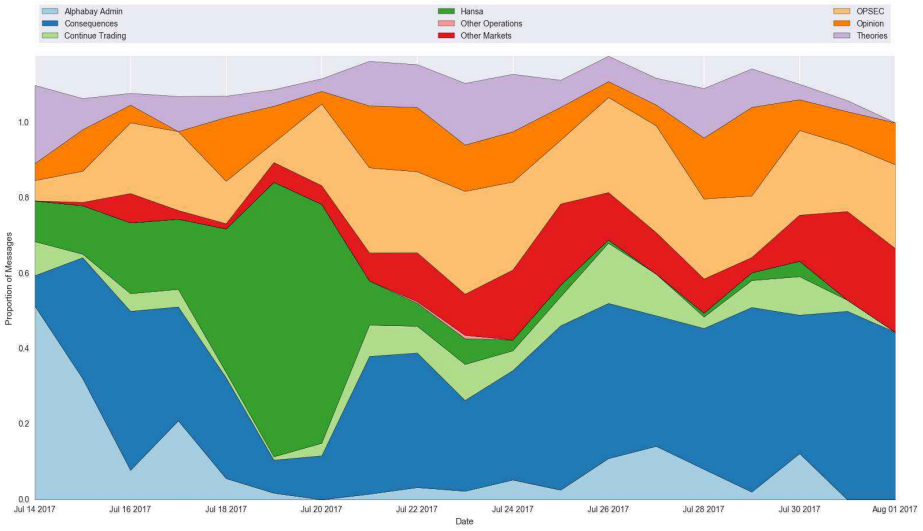

FIGURE 8.19: **Relative Proportions of Each Topic over Time**

- **OPSEC** (366 posts, 599 comments): posts and comments discussing how users can improve their OPSEC in the wake of the operation and recommendations on how users who felt vulnerable could protect themselves.
- **Other markets** (197 posts, 296 comments): posts and comments about active markets and whether or not they had also been affected by the operation.
- **Other operations** (9 posts, 9 comments): posts and comments about three other law enforcement operations which were mentioned in the discussion of Operation Bayonet, these were Operation Onymous, Operation Hyperion and Operation Titanium (Tools for the Investigation of Transactions in Underground Markets) a three-year EU funded project to identify common characteristics in online criminal transactions (European Commission (2019)).
- **Theories** (215 posts, 317 comments): posts and comments in which contributors share theories about how the operation was conducted, these include theories on how law enforcement were able to access the DNMs and for how long they were under law enforcement control.

Figure 8.19 shows how the the proportion of comments in each topic changes over time[2]. *Hansa* was most discussed in the days surrounding its closure and then became a much less dominant topic, as this topic was defined as comments about *Hansa* when it was still considered active. The proportion of comments discussing other DNMs increased after *Hansa* closed, which would be expected as *Hansa* was considered the best alternative after *Alphabay's* closure. The proportion of comments discussing whether or not it is

[2]As comments can be coded into more than one topic, the proportions can sum to a value greater than one.

safe to trade appears to increase marginally after *Hansa* shuts down, this could reflect greater instability within the ecosystem community.

Each of the themes and their content are now discussed in greater detail.

Alphabay Admin

The discussion of alleged *Alphabay* administrator Alexandre Cazes continued throughout the observation period. Contributors shared news stories, posted tributes, and speculated about the extent of Cazes' involvement in the running of the site. They also discussed his death and shared theories on how he was identified. Contributors posted 40 comments containing links to, or quotes from, news articles about Cazes' arrest and subsequent death. Information was primarily sourced from major news sites such as `www.wsj.com`, `www.independent.co.uk`, and `www.washingtonpost.com`, though two comments linked to `www.justice.gov`. In addition, contributors independently researched Cazes, uncovering his contributions to `www.rooshvforum.com`, the forum owned by Roosh V[3].

The information shared, and subsequent discussions, present two images of Alexandre Cazes. The first is of a talented programmer whose contributions to the ecosystem were appreciated. Many of the comments (29) along this vain expressed remorse at his death.

> He was an extremely intelligent guy. Being young, knowing this much about computer stuff and programming and all his abilities. Amazing. No one on Reddit could do even 10% of what this guy was capable of. It's sad it went wrong way [sic]

The other image is of a misogynistic and irresponsible person who tipped of law enforcement with his excessive and boastful lifestyle.

> He could've easily lived on the down-low in Thailand and never been caught. WTF is with the extravagant lifestyle when you have a chance of getting caught.

At least 47 comments argued that Cazes was easy for police to track down because of the way he spent his money (*"the seizure documentation talked about how he liked to brag about his wealth on the rooshv forum"*) and a further 13 argued that he made an error in living in Thailand because this enabled US law enforcement to access him (*"Thailand is not and never has been safe. That's why internet criminals avoid it and go to places like Brazil, Colombia, Costa Rica ..."*).

[3]Roosh V is a blogger and pick up artist who self-publishes sex advice books and travel guides

In addition to discussing Cazes as a person, his role as an *Alphabay* administrator was also discussed. Within the press, Cazes was named as admin DeSnake and as admin Alpha02. However some comments argued that he was not DeSnake (4) and/or Alpha02 (4). Though this was less than the number of comments (11) arguing that he was Alpha02. In addition, contributors disagreed as to whether he was the only admin or if there were others who had managed to escape custody. Some contributors argued that Cazes was, in fact, a "fall guy" for the other admin which was why he was the only one caught (*"That man was a set up while the rest of the admins skated with the coin"*).

Finally, contributors discussed the nature of Cazes' death, which was reported as a suicide. Some comments (31) argued that he was not dead but had been put into witness protection or bribed his way out of prison, or that he had been murdered in custody. These comments were in the minority compared to the comments that explicitly argued against such conspiracy theories or implicitly accepted the narrative found in the major news articles shared on the forum. However, they are part of a wider theme found in multiple topics discussed in the wake of Operation Bayonet which argue that law enforcement would deliberately spread misinformation and, therefore, not everything they say can be trusted.

> Too bad we'll never get his side of the story. Starting to sound like parallel construction and then the [sic] snuffed him out before he could leave clues about how he was actually caught.

This theme could weaken the deterrence impacts of arrests and other such interventions by allowing users who only learn of the event via conventional media sources to convince themselves it has not occurred.

Alexandre Cazes' death was widely discussed and left the forums divided in their opinion of him. This differs to, for example, the arrest of Ross Ulbricht who more conclusively received, and still receives support from the ecosystem (The Ross Ulbricht Legal Defense Effort (2018)).

Consequences

The consequences of the operations was one of the most discussed topics. 430 comments either described consequences users had experienced or speculated on future potential consequences and a further 116 queried the forum about the potential consequences. Contributors detailed many different consequences of the operations including the acquisition of data by law enforcement (89 comments); compromised accounts (123); losing access to the ecosystem (22), money (64), an order (62) or a trading partner (52); a loss

in convenience when trading (77); harm to drug users (4); and, impacts on related technologies (10). Though the specifics and the extent of these consequences were debated.

Because the Dutch police demonstrated they had access to *Hansa's* servers, contributors were aware that data stored on these servers was also accessible to law enforcement. As such, the form of this data and the implications of this were discussed. Some contributors argued that law enforcement had collected thousands of addresses ("*Feds in dutch used hansa as a honeypot for 1 month and manages to get 10,000 unencrypted addresses(world wide)*"), or even more information ("*LE now has your identity after using hansa ...*") though some contributors felt that only Bitcoin transactions were visible ("*If you PGP'd the address on your own machine all they'll have is the bitcoin transaction.*").

It was less clear if, in Operation Bayonet, the servers of *Alphabay* were seized. Some contributors argued they weren't ("*They never actually got access to the AB servers whilst the market was online*"), some argued that it didn't matter if they had because *Alphabay* automatically deleted messages after 30 days and so the servers did not contain much information. However, some contributors argued that law enforcement had acquired significant amounts of information and, therefore, former users should take precautions ("*Act as if they have the messages, because there is a high probability that they do.*").

The capture of data was of concern to contributors because of its potential to enable law enforcement to identify and arrest users. Some of the discussions covered what would happen to the users whose addresses had been collected. Though some contributors were convinced that there would be no consequences ("*even if they do raid you, the worst you're looking at is a caution*") others warned the operations would result in arrests, raids and busts ("*Investigations will be on going. Tons of sellers will get nabbed*"). Indeed two contributors claimed that the acquisition of data was more damaging for users than the loss of coin seized within the raid ("*F*** the money. You have bigger issues now.*").

The consequences of law enforcementobtaining data on ecosystem users is a much more dominant theme within the topic of consequences after 13 July. This is likely due to the fact that it had become known that law enforcement had access to *Hansa*, and potentially *Alphabay's*, servers.

Contributors also argued that, when *Hansa* was under law enforcement control, law enforcement were able to collect usernames and passwords for vendor accounts. These credentials were then used to log into vendor accounts on other DNMs where vendors reused their passwords. Indeed, some vendors reported being locked out of their

accounts. Because of the uncertainty of the multi-market closure, this led some contributors to speculate if other DNMs (particularly *Dream*) had also been compromised. However, the following quote from Dutch police was shared by a contributor:

> We have modified the source code which allowed us to capture cleartext passwords, PGP-encrypted order information, Bitcoins, IP-addresses and other relevant information that may help law enforcement agencies worldwide to identify users of this marketplace

And other, similar, confirmations countered arguments that law enforcement were in control of a third site.

Contributors used Reddit to publicise which vendor accounts had been compromised and circulate a list to warn buyers. One contributor claimed that at least 16 *Dream* accounts had been compromised ("*Dutch LE has at least 16 vendor accounts*").

Contributors also documented what they had lost during the law enforcement action. This included access to trading partners, accounts and losing products either because they thought they had been seized or because vendors had not shipped them under the cover the DNMs closing. The loss of Bitcoin was also documented, however this was to a smaller degree than in the wake of the *Alphabay* closure. This could be because contributors were more concerned with other consequences now they understood the role of law enforcement in the market closures or it could be because *Hansa's* system, which did not require customers to transfer coin into a wallet before making a purchase, reduced the amount of coin that was seized.

Another loss that emerged was a loss of convenience when trading on the ecosystem. Contributors observed a slower service ("*Some vendors are backed up on orders by of the AB crash.*") potentially due to remaining DNMs being overloaded by former *Alphabay* and *Hansa* users. Additionally, contributors described losing contact with trading partners and losing access to favoured products, some also felt that there were less products available on the remaining DNMs ("*There's less choice...*","*I'm in DHL but that does not have the s*** load of product that was on the other markets*").

Within this theme, some observations of how the ecosystem may have changed were also documented. Some contributors argued that shipping times were increasing ("*shipping times are delayed due to the AB shutdown and other things*"), that in transit seizures were increasing ("*They're probably deliveries intercepted because LE supposedly have 10,000 addresses from Hansa*") and that vendors were increasing prices and/or changing their refund policies as a way of recovering coin lost in the intervention:

> He could have been charging this high the whole time. He is clearly trying to recoup what was likely a substantial loss. In addition other products of his, which were not as hard to find, were also jacked up to higher than some competition.

Further, contributors linked phishing attacks and exit scams to the operations as, when DNMs go down attackers can create fake versions to steal credentials from unsuspecting users. Additionally, it was claimed that vendors were more like to exit scam during this period ("*there's also another [vendor] who is in a little bit of a trouble and can't wait to ramp up sales via FE to exit scam*") and other users may adopt missing vendors usernames and use their reputation to scam buyers ("*Every time a market falls people rush to DD because they feel it will be safer and every time the scams go through the roof*"). Several comments (56) documented a climate of wariness and uncertainty or an otherwise loss in confidence in trading on the ecosystem.

> Seeing the community drowning in mass panic and confusion makes me think [law enforcement are] playing pretty well, actually.

Some users were concerned about how the ecosystem in turmoil might affect addicted users or users who were dependent on the ecosystem for access to their products ("*Many addicts and dealers will be in some serious pain from this*").

Finally, some discussion was had about whether or not the double-site closure had impacted the value of Bitcoin. Some argued that the removal of so many coins at the same time had caused the value to crash ("*It dropped when AB shut down and I don't think it was a coincidence*") however this was disputed by 20% of the comments on this theme ("*I think you under [sic] estimate how little the DNMs effect bitcoin prices. It's the other way around... btc effects the DNMs.*").

Not all contributors felt that all users were at risk of repercussions. Just as, in the wake of *Alphabay's* closure, some contributors were chastised for leaving their coin on DNMs ("*Its always people who leave money in markets that get f***ed*") and therefore blamed for their loss, many contributors predicted the impacted users would be limited to those with poor OPSEC. It was argued that law enforcement would only be able to obtain the addresses of users who relied on *Hansa's* built-in encryption system and, therefore, users who had manually PGP encrypted messages would be protected ("*I'm sure there are gonna be pretty s***ty consequences to come for people who have made OPSEC mistakes. I'm fairly confident as I only ever encrypted myself*").

Contributors also argued that law enforcement were only interested in vendors and buyers making large, likely wholesale, purchases and that, therefore, small time buyers buying for personal use would likely be ignored. Additional attributes of interest included the products being purchased with some contributors claiming that fentanyl, opioid, and firearm buyers would be targeted.

> But they, more likely than not, will come after you if you're a vendor, bulk-buyer, and unfortunately I'd be inclined to believe if you purchased a somewhat large amount of opioids

A reason given for why law enforcement would focus on these users specifically was that they were under resourced and that it would take too long to investigate every user and so they would have to target priority users only. This was not an opinion held by everyone though, some felt that all *Hansa* and *Alphabay* users were at risk and some especially felt that, in this operation, law enforcement might make examples of small time users in order to deter others.

> I have a feeling that if LE are really smart then now would be the time to take action against a few personal buyers

One further caveat discussed was the location of users. Because the closure of *Hansa* was known to be led by Dutch police, two contributors reasoned that only European users would be targeted, however three others thought that it was US users who were most at risk. i.e., the discussion of which locations made users vulnerable to law enforcement appears to be led by the key law enforcement groups involved in each operation.

There was also speculation of when contributors might expect to see the repercussions of the operation, with previous operations being cited as evidence that law enforcement takes a long time to process information and conclude investigations. Some contributors speculated that the intervention was still ongoing and that more DNMs would be revealed to be compromised ("*You have to assume now, that any new market place is a law enforcement honey trap.*"). However, others argued that the operation had already been concluded and that any repercussions had already been felt ("*all safe, dont [sic] think any vendor has been busted due to the hansa and ab take downs*","*Glad we both made it. Lesson learned.*").

Overall, by the conclusion of Operation Bayonet and the closure of *Hansa*, some direct harms were recorded on the Reddit forums. Contributors described being inconvenienced by the intervention though fewer contributors explicitly recorded losing money when *Hansa* closed. The revelation of the role of law enforcement seems to be related to an

increased concern in further repercussions, for example arrests, as would be expected. The fact that the two markets were closed in quick succession also may have compounded the impact with some users reporting their losses across both events.

> I'm laughing but I really want to cry, I just went through the process of burning my drop and it's relative BTC wallets and all the usual because of Alphabay. Now my new one is a burn too

Despite this, many users seemed to think that this intervention would not seriously harm them or the ecosystem. They argued that most of the users would emerge unscathed and that this was merely a scare tactic which would not be followed up by action.

> It's basically just the terrorism approach. Get everyone to freak out and go into disarray by saying scary stuff like 'we've been collecting names, we're coming for you' etc. Realistically they don't have the time, money or care to go after you.

Some users were confident that the operations would positively impact the ecosystem. They considered operations such as this one to make users and markets more resilient and improve OPSEC by learning from which mistakes had been exploited.

> Remaining markets, vendors and buyers will tighten their opsec. Unless they have a tor exploit it will get much more difficult for LE to take down dnm's

Out of 77 comments that explicitly discussed the severity of the consequences, 66 claimed that the consequences were small or temporary or similarly fine whereas just 11 claimed that the consequences of these operations had actually been severe. This could be a result of the narrative constructed after both the closure of *Alphabay* and *Hansa* where by only the users who did not adequately protect themselves were considered to be at risk.

Continue Trading

A key theme that emerged was whether or not contributors would continue to trade on the ecosystem. In 189 comments, contributors explicitly referred to the decision of returning to trade after the closure of *Hansa*. In 102 comments, contributors declared they would continue to trade. Some of these decisions were caveated, however, as some contributors stated they would only trade through direct deals ("*if you absolutely must purchase, do direct deals with trusted vendors*") and others that they would only make

small personal orders ("*if I can't hold back I'll only go small incase [sic] of an exit scam or LE takeover*").

A further 12 comments were queries about whether or not it was safe to trade again and, if so, which market they should create an account on, implying that the contributor had the intention to continue trading though had not done so yet.

By comparison, in just 32 comments contributors claimed that they would stop trading. Some argued that it was no longer safe to trade; others that there were no opportunities to trade, i.e. no viable sites or vendors (" *as far as anyone around here should be concerned there IS NO DNM right now*"); and some claimed to have a sufficient supply that they did not need to use the ecosystem ("*I ain't risking it myself though, I have a stockpile to keep me going*"). One contributor argued that many users were taking time away from the ecosystem in the wake of the operations.

> Look at the numbers - Dream 13,000 registered users, Traderoute - 7000 registered users, Wall Street - 1000 registered users.AlphaBay had 200,000 active registered users

In 58 comments, the contributor demonstrated an intention to continue trading but not at the present time. For some users, this meant waiting a few weeks ("*Yeah don't be stupid, I wouldnt order anything for a week or two but its not the end of the world*") or months ("*...we feel safe placing orders in 6 months time*") and for others that they would wait until they felt the ecosystem was safer or no longer monitored by law enforcement ("*I'll be registering on there once things cool down.*").

A discussion was also had about whether or not contributors would start or restart trading offline. Some argued that this was their option now that they could not use the ecosystem ("*They've just made me spend more money for worse quality in the streets for a little while*") however others felt that trading online was still preferable.

This shows that, after *Alphabay* was revealed to be a law enforcement operation and *Hansa* was closed, a smaller proportion of contributors made claims that they would continue trading, even though they were still in the majority. Further, there were more claims from users that they would only trade through direct deals.

When some contributors expressed intentions to continue trading they did so by naming the DNMs they intended to create accounts on. *CGMC*, *DHL*, *Dream*, *Sourcery*, *Tochka*, *Traderoute*, *Valhalla*, *Wallstreet*, and *Zion* were all named as viable alternatives. Some argued that *Dream* was not safe ("*Personally I'm avoiding dream until we know what's what.*") or that it was not sensible to trade on an existing market ("*Yea I would advice [sic] you to wait for a new market to surface tbh*").

This discussion is very similar to that observed at the start of the measurement period, after *Alphabay* first went offline. However, contributors were now concerned that *Dream* may be compromised, as opposed to just wary it would exit scam.

Hansa

When the official press release on Operation Bayonet was published (on 20 July 2017), contributors began sharing news stories about *Hansa's* closure. Their sources included krebsonsecurity.com, motherboard.vice.com, justice.gov and the Europol press release and Attorney General's press conference directly. The vast majority of these 40 or so comments just shared links to or quotes from news stories, without sharing any opinion or emotion. Though, those that did, most often shared exclamations of surprise or fear ("*oh s****", "*f****", " *Unlucky dude. I would head for Belize*").

Prior to *Hansa's* closure, the conversations that centred around the marketplace discussed its merits and flaws, how vulnerable it was to closing and how to use the site. In addition, a key topic to emerge was the decision of the marketplace to ban the product fentanyl and whether or not this action was shared by the community.

34 of the comments described *Hansa* positively arguing that it was the best alternative now that *Alphabay* was closed and citing reasons such as the good interface ("*As for Hansa, it's just the interface and UI alone that are a selling point for me*"), the lower number of scams ("*New markets are filled by scammers, dream market is too, only hansa has any decent vendors...*"), and its security ("*Most of us probably enjoy the security a market like hansa offers*").

Some contributors also used the fact that *Hansa* offered multisig transactions to claim that it was protected against an exit scam ("*Multisig factors out a market walk*").

In contrast, 24 of the comments described *Hansa* negatively. Contributors considered *Hansa* to be expensive ("*Hansa man Hansa... Transaction fees tryna bleed us dry*") and lacking in variety and availability for both products and vendors ("*they still do not seem to have as much listings and variety than Dream*").

During the short period before *Hansa* was closed, it went temporarily offline causing contributors to speculate that it would exit scam ("*Then hamza [sic] will either crack under the pressure of all the refugees, or they'll find a way to exit scam too.*"). However, very few comments could be found which actually predicted *Hansa* would also be closed by law enforcement, implying that this element of the operation was not anticipated.

On 20 July 2017, just before *Hansa* was closed, some users posted about withdrawal issues that they were having. This was the first substantial indication that users had of the site's impending closure however none of the comments that complain of this

issue speculate that the site would close (despite how, during the discussion of what had happened to *Alphabay*, contributors identified coin withdrawal issues as a first indicator of an exit scam).

On 18 July 2017, *Hansa* administrators banned fentanyl from the site. This decision was met with a debate about its legitimacy with 47 comments arguing that banning any kind of drug was in antithesis of the Libertarian values of the ecosystem which supported fully free trade ("*The whole DNM culture is built around the idea that you can't restrict a market when there is a demand– and there IS a demand for fentanyl*"). Similarly, it was argued that it was hypocritical to ban this particular drug as opposed to any other:

> lol. people sell all sorts of dangerous s*** on the markets, but oh boy lets [sic] get on our high horses and act like we are morally good people by banning fent because its [sic] dangerous, but heroin is a-okay. what are they going to ban next?

However, nearly as many (46) comments agreed with the decision, claiming that fentanyl was not only too dangerous to take but also attracted law enforcement more than other products. Some users claimed that DNMs that did not allow the sale of fentanyl would be less likely to be targeted and shut down ("*fent is gonna be the downfall of dnm's ban the s***.*").

This argument was not believed by all users, however, with some comments claiming that banning fentanyl would make the drug more dangerous because vendors would no longer declare when they had pressed it into pills or used it as a cut with another substance. The decision was divisive but, ultimately, has been replicated on other markets, for example the *Dream* marketplace (C.M. (2018)).

Whether or not this action was driven by law enforcement or the actual *Hansa* admin was speculated upon by users after information about the whole operation was released. Some users felt that the idea had come from law enforcement because they were not allowed to enable fentanyl sales whereas others cited interviews in which Dutch law enforcement imply that the decision actually came from the administrators themselves (Krebs (2017)).

Opinion

In 334 comments from 189 posts, contributors shared their opinions on, and reactions to, Operation Bayonet and the closure of *Hansa*. More specifically, they shared their opinions on whether or not the operation had been successful and if they considered it to be a legitimate or ethical operation.

In 35 comments across 20 posts, contributor explicitly named emotions they were feeling, for example that they were "gutted" or that the incident was "irritating" or "depressing". In 20 of the comments (from 9 posts) the emotion expressed was negative, however in 15 comments (across 12 posts) users expressed admiration for the operation itself, often begrudgingly.

> Honestly, I'm just impressed with the whole honeypot thing. Like, I'm mad, but *god damn* that was well orchestrated.

112 comments across 78 posts discussed whether or not the operation had been successful, i.e. the contributor made a claim as to whether or not the market closures would stop future trade on the Dark Web. In the majority of cases (92 comments, 65 posts) the contributor argued that the operation had been unsuccessful. These contributors argued that law enforcement had not made the achievements they claimed ("*Where all [sic] all of the arrests if they got all this info.*"), that the operation had not gone to plan ("*the "we made alphabay look like an exit scam" line is post facto FUD from the FBI - tehy wouldave [sic] loved to ave [sic] run that market for a few weeks or months and looked a bit more pro like the dutch [sic]*"), that this operation would actually have a negative impact on current investigations ("*In reality, Alphabay shutting down essentially ends all open investigations sending these agencies so far back*") and that this operation, as with any other, is ultimately futile because of the resilient nature of the ecosystem:

> Every time we have a massive blow to our operations, our institutions and our shared wisdom, we never give up, we come back better and stronger. Look at the markets in 4 weeks, 8 weeks and 6 months time. We'll probably be bigger than ever.

In contrast, just 21 comments from 19 posts argued that the event had successfully damaged the ecosystem. Some contributors made observations about users that they already felt had been affected ("*You just have to look at some other DNM sub reddits to see how hysterical some people got*") and some claimed that this particular operation was the worst one they had seen ("*I have been a part of the community since SR1 and this bust has been the worst one yet*"). Some contributors expressed the idea that this operation would fundamentally change the way the ecosystem operates by removing such a large and dominant marketplace ("*I think the time of huge market places is over*").

However, even comments that argued this also often presented this change as a surmountable one which did not necessarily result in an entire ecosystem collapse. For example, one such comment argues that law enforcement had changed their tactics to

become less predictable and more effective, that they had "*gained tremendous ground*" and that, as a result of the novelty of the ecosystem wearing off and the DNMs become too unreliable to support users with drug habits, there were fewer and fewer users: "*I don't see a lot of new people coming into the community and I see a lot of staples either being taken down, easing into the background to conduct business, or scamming*". However, the comment still concludes:

> Is it worth it?... My answer to that a few years ago was a resounding "Yes!!!" My answer now would be more along the lines of, "meh, I need to learn all that s***, and it's a bunch of scams anyway." What is the community planning to do about this?

I.e. even users who consider the ecosystem to be weak and vulnerable also believe it can be redeemed and improved.

Further, 55 comments (on 42 posts) called other contributors out for spreading FUD. For example:

> *Sound* speculation is indeed beneficial. However, in line with the OP, far too many people (quite obviously) either don't possess, or choose to discount, some very basic principles of logic. There's also an overwhelming lack of understanding of the finer points of *much* of the technology involved. *That* kind of speculation = FUD, which isn't helpful to anyone.

The high presence of FUD comments implies that the negative speculation about the future of the market was not accepted by all other contributors and was a concept actively fought against by some.

Finally, in 70 comments across 47 posts, contributors discussed whether or not they considered the action taken by law enforcement to be illegitimate. As with the discussion of success, most users felt that the operation was not legitimate. Indeed, in 62 comments on 43 posts, this sentiment was expressed vs just 8 comments on 5 posts in which contributors argued that it was a legitimate legal action.

Contributors argued that Operation Bayonet and the *Hansa* honeypot were illegitimate because they facilitated the sale of drugs ("*Funny thing is though haven't LE actually gone against the law by hosting hansa on their servers and allowing transactions to take place*") but also because they felt that targeting online markets would make the drugs trade more dangerous for consumers ("*I guess LEO want to see a rise in crime due to people doing IRL exchanges for drugs*").

Additionally, some contributors argued that the criminalisation of drugs was itself hypocritical and illegitimate:

> Ah well at lest the government can keep selling us booze and cigs, legal drugs are best drugs eh, f***ing hypocrites

When contributors learned of the law enforcement operations that closed *Alphabay* and *Hansa*, they evaluated their success. Success was judged on how effective the operations were at challenging the ecosystem and protecting drug users. Whilst the first objective was disputed, the consensus on the forum was that targeting DNMs would make it more dangerous to buy and use drugs.

OPSEC

There were 278 comments across 186 posts in which contributors made recommendations or dispensed advice on how to improve OPSEC in the wake of the operations. There were a further 29 comments across 21 posts in which contributors asked queries about how best to improve their OPSEC. There were lots of different recommendations which covered how to prepare for an investigation and how to keep trading securely.

The most common piece of advice for users who were worried they might be investigated was to clean house, i.e. remove any traces of drug activity from their home so law enforcement would not have physical evidence. It was also recommended that users cleaned their online presence, for example by deleting their account, changing their username and passwords, or by destroying their hard drive etc.

> If they don't physically find drugs or a computer used to make any order, I don't see any realistic way they'd be able to charge a buyer. Unlikely they could even charge them unless they actually caught you with drugs or caught shipments. So clean house of drugs. If you can also get rid of your computer, do the same.. Lay low for a month or two, keep watching. See whatever emerges next.

This differs to the discussions around OPSEC after the initial closure of *Alphabay* which was more focused on protecting a user's assets and primarily about actions taken online.

Some of the advice was contradictory with some contributors recommending that users changed their PGP keys and others advising against this as it made them look less trustworthy ("*you should never have changed your public key. this is the only thing to prove you are who you say you are.*"). Evidently, erasing a user's online history to protect themselves with law enforcement has an impact on their future interactions in the ecosystem.

The second most common recommendation for users on how to protect themselves was to lay low. Some contributors argued that users need only do this for a short while until law enforcement were no longer able to continue their investigation ("*Wouldn't those warrants also go "stale"? Meaning if those vendors lay low for a couple weeks they're clear*").

Despite the number of contributors advising others to lay low, or perhaps because of this being characterised as a temporary measure, many users were seeking advice on how to continue to trade more securely. For these users, a number of technologies were recommended as more secure, for example I2P as an alternative to Tor, Tails, and VPNs. As some users were concerned that law enforcement were tracing Bitcoin transactions to identify users, tumbling and Bitcoin ATM's were recommended as methods of extracting coin without leaving a trace.

In addition, alternatives to Bitcoin that were considered more secure were discussed with users recommending Dash, Ethereum, Monero, and Zcash. Monero was overwhelming the most recommended cryptocurrency (receiving nearly 50 recommendations).

The most common recommendation for users who were going to continue trading was to personally encrypt their messages using PGP encryption, rather than relying on auto-encrypt protocols. This was part of a wider theme of recommendations in which contributors argued that, to stay safe, users simply needed to maintain good OPSEC.

> Everything/everyone is potentially law-enforcement,, or other evil-doer, trust no one, so pay close attention to your s***... the darknet is not safe, don't go there if you're a child. Be an actual adult, learn about what you're doing, take precautions, know how effective your precautions are, so you're in control of how "dangerous" your activities are. Then when these inevitable crises occur, you won't be freaking out about how much of your ass was/wasn't hanging out in the breeze

However, because of the specific nature of the *Hansa* bust, some discussions accepted that good OPSEC practices were not completely in the control of individual users. The fact that some of the vendors on *Hansa* had their accounts compromised on other marketplaces meant that users had to ensure that they also weren't communicating with compromised vendors ("*Buyers just need to avoid placing orders with these vendors until they're banned on Dream or support gives access back to the actual vendors*").

Additionally, contributors discussed what kind of security new markets should ensure. This discussion occurred after the closure of *Hansa*, rather than *Alphabay*, perhaps because it was sparked by the knowledge that law enforcement had closed two major

DNMs. As such, when contributors became aware of vulnerabilities in the ecosystem structures, rather than just individual accounts, this created a need to improve more than just their own OPSEC.

Desired improvements in market OPSEC included having a built-in auto deletion of messages, no auto-encryption and for the servers to be hosted in countries with less active law enforcement. These are all elements that contributors also argued made either markets or individuals vulnerable in the closure of *Alphabay* and *Hansa*. Some contributors also called for future markets to ban fentanyl and fraud which were felt to be reasons why these particular markets were targeted.

Other desirable security features were for any future market to be decentralised and for it to facilitate multisig transactions. The conversation around multisig transactions explored both whether it is a positive feature or not. Some users argued that it was difficult or too complicated to use ("*I think the multisig is too much work*"), made users more vulnerable to being identified through their transactions ("*hansa multisig is going to get lots of vendors busted... The real inbound and outbound btc addresses are obtainable more easily than an autotumbled singlesig transaction.*"), and offered no protection for coins being traded on a market, as evidenced by *Hansa*.

However, those who defended multisig argued that it should be on future markets because it did mitigate the harms of an exit scam or market seizure, especially before *Hansa* closed ("*Come over to Hansa. Multi-sig is the way to go if you don't want to get f***ed out of your money.*"). These users also argued that *Hansa* simply did not use a true implementation of multisig and so the issue was specific to that site and not the technology itself.

Many comments (108) raised questions about how *Hansa* users could recover their coin after the market was closed. The *Hansa* multisig set up was designed so that, if the market ever closed, after 90 days any coin held in market wallets would be released. Some contributors argued that Dutch law enforcement had changed the implementation so that any of these coins would be paid into wallets controlled by law enforcement when they closed the site ("*after LE took over the market they changed the source code and now no one is getting any money from either escrow or multisig*") but this theory was disputed ("*I verified the multisig transactions, they're valid 2-of-3 transactions partially signed by Hansa (valid signature). Output address is still the one I personally control.*").

The discussions on OPSEC after the closure of *Hansa* compliment those in the wake of *Alphabay's* disappearance. There is a strong theme of individual users taking responsibility for their own security and this being seen as sufficient protection to continue

interacting with the ecosystem. However, after new information on the role of law enforcement in each operation emerged, new conversations about how the markets themselves could be made more secure were had. These conversations appear to be directed by how contributors felt that law enforcement were able to close down each market and, where information was missing (for example in the case of if users would recover their funds from the *Hansa* multisig implementation), the correct course of action was disputed.

Other Markets

There were 295 comments across 195 posts that discussed other markets still active after *Hansa* closed. These markets were *DHL*, *Dream*, *Greenroad*, *Sourcery*, *TMG*, *Tochka*, *Traderoute*, *Valhalla*, *Wallstreet*, and *Zion*. *Dream* was the most talked about alternative market and was mentioned in 237 comments, the next most mentioned comment was *Traderoute* which was discussed in 25 comments.

Other markets were discussed as viable alternatives and in the context of the operation. *Greenroad* and *Tochka* were recommended, as was *Zion* but it was observed as having no US vendors ("*Zion looks pretty good, but yeah, no US vendors right now is a problem*"). *TMG* and *Valhalla* raised suspicion by closing registration and having withdrawal issues respectively, as these behaviours were considered warning signs that each market had been compromised ("*We should assume that all markets that close registration are under control by LE, I'm going from past experience.*"). The discussions about *Wallstreet* and *Sourcery* were about how they had also closed during the measurement period.

For the larger markets (*DHL*, *Traderoute*, and *Dream*) their viability and security were more debated. Nearly two thirds of the comments discussing *DHL* recommended it, primarily because it was an invite only market and therefore more secure:

> You're completely missing the point of DHL. It's only the trust/elite/vetted vendors on there, not every average Joe. The whole point is you have a safer and more exclusive environment while also building a relationship with top-tier vendors.

However, some contributors argued that it was a honeypot or would be next to close. A piece of evidence for this was that it was being highly recommended on the forum. Users argued that, as *Hansa* was recommended by lots of users and then closed law enforcement might have been influencing the movement of users from *Alphabay*. Therefore, users should avoid the next most recommended market ("*It was 100% not a coincidence that immediately after (and even before) AB went down that posts with Hansa were being upvoted. If you were smart that should have ended any notion of using that market. Hype over illegal activity should be an immediate red flag.*").

Similarly, approximately two thirds of the comments discussing *Traderoute* recommended its use. The reasons given for switching to *Traderoute* were that it was established, and therefore wasn't untested, that the support on the site were good ("*In particular I hear that the support is actually good, which is a welcome change from the incompetent idiots that ignore customers that I'm used to from places like AlphaBay and SR 2.0*") and that it was similar to *Hansa*. However, some contributors warned against *Traderoute*, again, because it was receiving so many positive reviews and this was an indication it was the next honeypot.

By contrast, just one quarter of the comments discussing *Dream* claimed that it was a safe or good market to begin trading on. *Dream* was, at the time, the largest DNM and had been operating for the longest but some contributors had issues with the administrators of the site and had concerns it was controlled by law enforcement. There were several complaints of technical issues and this, and the fact that it was reported to be offline several times during the observation period, fed into a rumour that *Dream* had a security flaw which would leave it exposed. As several vendors were also reporting their accounts had been compromised, several contributors argued that *Dream* was unsafe:

> it's my personal opinion dream market is rogue market ran by police or officials on the inside trying to steal money and gather intelligence.

Despite all of this, some contributors still felt that it was the best alternative DNM or, at least, the only available option for trade:

> It's the only market with lots of vendors and good prices. All other markets have a small amount of quality vendors however the prices are insane

A review of the comments discussing alternative markets to *Alphabay* and *Hansa* shows that, after *Hansa* closed, many contributors were unsure of which market to start trading on. Many factors were taken into consideration including the range of products and vendors available and the usability of the site. However, because of the specific element of the double operation in which the most popular alternative market (*Hansa*) was shut after *Alphabay*, the debate became further complicated with contributors left feeling unable to recommend the markets everyone else was recommending.

Other Operations

Operation Hyperion was mentioned 6 times in reference Operation Bayonet from 20 July 2017. 1 comment explained how a list of usernames being circulated as a consequence of Operation Bayonet was actually a photoshopped version of the list of known users from

Operation Hyperion and 1 comment was arguing that Dutch police were "*pretty hot on DN stuff*" as evidenced by their involvement in Operation Hyperion.

The remainder of the comments used Operation Hyperion to diminish the speculative consequences of the event. Contributors argued that, as with Operation Hyperion, the consequences of this event would be a warning letter but no action for low priority users.

> Chill out man unless you was [sic] buying kilos of bud weekly of the said vendor worse that will happen is letter threw [sic] door couple of years from now just like they did with Silkroad

Operation Titanium was mentioned 3 times and only between 21 July and 22 July 2017. Each comment hypothesised that the two operations were linked, but were unable to evidence the claim nor did they share thoughts on how this might affect Operation Bayonet.

Operation Onymous was mentioned in 5 comments, mostly as a point of reference within the history of ecosystem or to remind contributors that law enforcement operations occur periodically. The way in which this operation is discussed is as being bigger than Operation Bayonet, though this is not commented upon (for example to claim that this operation is, therefore, more impactful). One contributor advised other vendors to set up direct deals with their customers, an action that they took after Operation Onymous.

> After Operation Onymous, I was contacted by a few of the customers I had dealt with on C9, and we've been doing direct deals for over 2.5 years now.

As with the discussions after Operation Hyperion, forum contributors demonstrate a knowledge of other operations and can use these to cast judgements on operations currently active. Therefore, operations do not have only an immediate impact but can help to increase, or decrease, the impacts of future operations.

Theories

There were 292 comments across 193 posts in which contributors discussed theories about how law enforcement were able to close down *Alphabay* and *Hansa* and their behaviour during the operations. The theories that arose were that the administrators of the sites identified themselves or the locations of the servers through some form of error or by going to the police, that law enforcement had exploited a security vulnerability in either site or in Tor, or that they had been able to trace Bitcoin transactions linked to either marketplace.

The most dominant theory was that site administrators were responsible for the closure of each site. 85 comments argued that this was the case. Some contributors simply expressed that admin were careless ("*AB, Hansa and Silk Road all got caught because they were extremely sloppy*") however, statements made by law enforcement informed more detailed theories. It was claimed that Alexandre Cazes was identified because he left his personal e-mail address in the header of welcome e-mails send to *Alphabay* forum users and this led to the location of *Alphabay's* servers

> It was the software in which the Alpha Bay forums were powered by (XenForo I believe they were?). It seems evident now that probably for the first day, maybe a bit more maybe a bit less, the software was configured to send out e-mails with a from: with his address in the headers.

For *Hansa*, there were two competing theories. One was that a site admin was arrested by law enforcement found child pornography on his work computer leading to an investigation,

> He downloaded in work childporn. They knew it was downloaded to that company and the forensic searched the computers of the employees for who it was and found the Hansa related stuff

And the other was that the admin were identified through an investigation into a clearnet site for e-books

> the hansa explanation i side with is the one where the admins also owned a clearnet site that sold ebooks illegally, and they raided that and also found the hansa site coincidentally.

However, these theories were disputed by other contributors. 12 comments argued that law enforcement did not exploit admin error or specifically refuted the e-mail theory (mostly by arguing that they had never received a welcome e-mail). A further 42 comments discussed the possibility that law enforcement were spreading misinformation as part of their operation. This was often referred to as "*parallel construction*".

> This part still does not make sense to me and reeks of parallel construction. No-one else ever saw that email.

Contributors expressed that law enforcement had been able to shut down *Alphabay* and *Hansa* through other methods they did not wish to reveal and created plausible stories

to hide this fact. A motivation for using parallel construction, given by contributors, was that law enforcement intended to use this method again. This, and the fact that two DNMs were shut, led some to believe that there was a vulnerability in Tor enabling the operations.

> Them getting busted one after the other says otherwise, they wouldn't locate a server and then wait until they can do the same for another before taking the first one down. This has Tor exploit written all over it.

The Tor vulnerability theory was advocated in 34 comments and only actively disputed in two. However, in a further 16 comments contributors argued the vulnerability was in the sites themselves (and therefore not in Tor). Several contributors argued that law enforcement relied on other services, such as Bitdefender and Coinbin to identify this vulnerability. In either case, some contributors who felt that there was a common vulnerability between the two sites also expressed concern for the safety of those continuing to trade on the ecosystem

> It's been said before but stay the hell off all markets for a while. We don't yet know whether something - PGP, the network, a particular market, a vendor - is fundamentally compromised in a very serious way.

Contributors also shared theories about how law enforcement operated and the exact nature and objectives of the operation. Some contributors felt that the operations were still ongoing and others that they had been operating for longer than was being claimed, particularly that *Hansa* had been operated by law enforcement from the outset ("*Maybe the LE owned Hansa right from the start*").

In addition, contributors discussed if the two market closures were connected or if it was a coincidence. Some contributors who felt the closures were connected also felt that there would be more closures. In contrast, those arguing the operations were coincidentally timed were also warning other users to not become paranoid about this exact eventuality:

> How do you know that the page you linked isn't misinformation to make it seem as though there was more coordination than there actually was, and that they are in fact capitalizing on a *coincidence*? Police agencies make busts all the time and create their own versions of events in order to obscure illegal investigative techniques. (see parallel construction). In this case I think they made this s*** up to instill fear in the DNM user.

Several different theories were discussed by forum contributors on how both *Alphabay* and *Hansa* were closed. Whilst official statements were circulated and informed dominant theories, contributors were concerned about law enforcement intentionally spreading misinformation both because their lies could be protecting a greater threat to the ecosystem or because the operations had been aided by luck and they wanted to spread fear across the ecosystem's users.

Summary

The discussion after it was revealed that law enforcement were behind the closure of *Alphabay* and after the closure of *Hansa* covered similar topics to the discussion after *Alphabay* closed, however not all of the themes within those topics were similar.

Contributors continued to speculate on why and how *Alphabay* had closed and, additionally, had this same discussion about *Hansa*. However, they now had access to more information from official law enforcement statements. This led more contributors to argue that the sites were indeed closed by law enforcement and ended the discussion about whether *Alphabay's* closure was temporary or permanent.

The specific methods used in each operation were still disputed and many contributors were unwilling to accept that official statements were accurate and believed, instead, that law enforcement were sharing misinformation about each operation.

The topic of the consequences of each operation remained prominent. However, the conversation appears to shift from talking about economic losses to the potential of being arrested. This is likely due to greater clarity in the role of law enforcement in the site closures. It could also be because *Hansa* employed a multisig protocol and so less coin was lost in this operation.

The conversation about multisig became relevant after the closure of *Hansa*. Prior to this event may contributors recommended it but some felt that the closure of *Hansa* demonstrated its inefficacy. There was not consensus on this issue, however, as those defending multisig protocols argued that the issue was with the *Hansa* implementation, not the technology itself.

More comments after 13 July advised that users stop trading on the ecosystem, at least for a while, than before. This could be because contributors were more concerned by the presence of law enforcement than exit scams. Alternatively, it could be because contributors were worried that there would be more site closures after *Hansa* was also closed. In either instance, the announcement that *Alphabay* was shut down by law enforcement in a manner connected to the closure of *Hansa* may have increased the impact of the operation.

Related to this theme, advice from contributors on how users could protect themselves going forward shifted from just focusing on good online protocols (such as not storing coins on marketplaces) to offline actions as well (e.g. cleaning house). However, in both discussions, contributors argued that the users most in danger of repercussions were those who were not careful enough.

For the most part, the same DNMs that were recommended to users looking for new accounts before *Hansa's* closure were recommended afterwards, with the exception of *Hansa*. However, one way that the closure of *Hansa* affected this conversation was by making some contributors concerned that the DNM they were most likely to move to or recommend was the next one to close. This was a benefit of the staggered multi-market closure as it appears to have made some contributors less certain of how to continue trading.

Chapter 9

Discussion

In this chapter, the results of the Cross Market study, Silk Road 2.0 study, and Reddit Forums study will be discussed in the context of the hypotheses outlined in Chapter 3. Then the limitations of the research and how these inform the conclusions of the research will be explained. Additionally, new findings not predicted in Chapter 3 will be presented and future work to further develop these findings will be suggested. Finally, this material will be used to make recommendations for law enforcement interventions on the DNM ecosystem going forward.

Evidence was found to partially support both models presented in Chapter 3. In this chapter, the results of this work are used to argue that, if a users loses their account or their reputation they are less able to continue trading on the ecosystem. However, if they only lose money or are made to feel the ecosystem is unstable, their capacity to trade is not diminished. If they lose trading partners, a product (that they are either selling or purchasing) or are more aware of the presence of law enforcement then they may be able to continue trading at a reduced capacity.

It is also argued that, at an ecosystem level, reducing the overall trade or convenience of trade on the ecosystem affected its resilience. However, reducing the amount of money in the ecosystem or its population did not.

The limitations of each study are also discussed, with the key limitations being the quality and completeness of each dataset employed. The implications of each limitations on the findings are argued and these are used to propose future research. For example, it is explained that the discussions extracted from the forum data are not necessarily representative of the actions of users and so a quantitative study is needed to explore the material results of Operations Hyperion and Bayonet. Additionally, the dataset used

within the Cross Market Study is incomplete in several ways and so may have obscured ecosystem growth after events.

Finally, the results of this research are interpreted into recommendations for law enforcement. Based upon the analysis of vendor responses to Operation Onymous, it is argued that operations should be targeted at high priority users as large events seem to have a greater impact on potentially less important vendors. Further, the comparison between discussions about Operation Hyperion and Operation Bayonet imply that users are more affected by operations that actually impact them, rather than simply seek to deter them and that the impact of an operation needs to be immediate.

9.1 Hypotheses

Each of the hypotheses proposed in Chapter 3 was answered by at least one of the studies presented in this thesis. The conclusions made about each is now presented. First the hypotheses pertaining to user behaviour and resilience are discussed (hypotheses 5-19) and the user model in figure 3.2 is evaluated. Then the hypotheses (hypotheses 1-4) about the overall resilience of the ecosystem are discussed and figure 3.1 is evaluated.

9.1.1 User Resilience

To understand if a fear of instability diminished a user's capacity to trade, Hypothesis 5 was evaluated to see if users who had witnessed many DNMs close (i.e. experienced instability) were more likely to stop trading after an event.

However, the logistic regression model built to evaluate the user population after Operation Onymous showed that the more DNMs a vendor had witnessed close the more likely they were to continue trading. This means that experiencing greater stability in this form does not appear to diminish the capacity to trade of the users evaluated.

Despite this, some evidence was found within the comments of the subreddits used to evaluate Operations Hyperion and Bayonet that contributors were wary of using the ecosystem in the wake of a market closure because of scams. They were concerned about phishing links to fake DNMs and malicious actors impersonating their previous trading partners. Therefore, some evidence was found of a fear of instability, but it was not found through answering Hypothesis 5.

Events that caused users to lose accounts were predicted to diminish a user's capacity to trade in multiple ways. Most directly, Hypothesis 6 predicted that users that lost

accounts in events were less likely to continue trading. This was shown to be the case after Operation Onymous and the Evolution Exit Scam as both events most dramatically affected the populations who lost all their accounts and then the populations who lost at least one account but barely affected the populations who did not lose any accounts.

Hypothesis 7 argued that the more accounts a user has, the more likely they are to continue to trade as they are more resilient against this impact. However, analysis of the population after Operation Onymous shows that only those users who had multiple accounts under different usernames were more likely to continue trading and that the number of accounts owned by a user negatively impacted the probability that they would continue to trade. Therefore simply owning more accounts does not necessarily make users more resilient to events in which an account is lost.

Losing an account was hypothesised to impact on users in three further ways: by causing them to lose reputation, lose trading partners and lose coin. These losses are articulated in Hypotheses 8, 9, 10 respectively.

Evidence for each of these hypotheses is found within the comments of the subreddits used to evaluate Operations Hyperion and Bayonet as well as the closure of *Hansa*. However, the relationships were not found to be as straight forward as every user who lost an account also lost coin, reputation and/or their trading partners.

After *Alphabay* first went offline, 5 comments were found in which contributors claimed to have lost their feedback. However, there were also comments from users claiming that other sites had been able to verify their user status on *Alphabay* and had therefore transferred their reputation.

There were at least 54 comments in which contributors described losing a trading partner. This occurred, not precisely because they had lost an account but because they had lost the platform through which they communicated with their trading partner and/or because their trading partner did not have another account that allowed them to continue trading.

At least 174 comments reported losing coin because contributors had lost an account on *Alphabay* and *Hansa* with some users reporting significant sums in the tens of thousands of dollars. However, there were many other comments arguing that contributors would not lose coin if they did not store any on their account. Further, there were less reports of lost coin after the closure of *Hansa* than *Alphabay*. This implies that losing an account does not necessarily result in a loss of coin. It also appears to be an impact that more greatly affects vendors as these users receive coin from buyers and must rely on the DNM allowing them to empty their wallets whereas buyers only need to refrain from adding coin to their wallet.

Losing coin was hypothesised to diminish a user's capacity to trade regardless of if the user also lost their account. Hypothesis 11 argued that users were more likely to leave after an Exit Scam, in which users were unable to withdraw any coin they had on the site, and an amicable site closure, in which they were.

To test this hypothesis, the closure of the DNM *Darkbay* and the Evolution Exit Scam were compared. These events had approximately the same impact on the population trading during each event but was not trading on *Darkbay* or *Evolution*, however less *Darkbay* users continued to trade after this site's closure than after *Evolution*. This directly contradicts the hypothesis.

Darkbay was a much smaller DNM than *Evolution* and potentially therefore had a more loyal user base which would also explain this result. Further, *Darkbay* closed when it merged with another DNM, *Andromeda*. The data from *Andromeda* was not complete enough to include in the dataset and so it cannot be measured how many users from *Darkbay* actually continued to trade on this new platform.

It was also found that, even though the Evolution Exit Scam removed proportionately more coin per vendor than Operation Onymous, it still had a smaller impact on the affected population. These results could imply that, actually, losing coin in an event does not diminish a user's capacity to trade.

Additionally the impact of the hack of *Silk Road 2.0* on the site's population and the impact of the Evolution Exit Scam on its population were compared. This is because, in the hack, users only lost coin whereas in the exit scam they lost coin and accounts. This comparison was used to evaluate Hypothesis 12.

After the hack on *Silk Road 2.0*, 95% of UK users continued to trade (and 100% of UK vendors), a far greater proportion than the proportion of affected vendors who continued to trade after the Evolution Exit Scam. This is confirmed using the Chi-Square test as a significant result (statistic = 20,578, p–value = 0.0).

The amount lost in the hack is much smaller than the exit scam, however this result, especially when combined with the comparison of an exit scam and site closure, imply that losing coin is more likely to reduce the chances of users trading when combined with losing an account.

If an event causes a user to lose reputation, either directly or indirectly, this is predicted to diminish a user's capacity to trade. Hypothesis 13 argued that users with lower reputations are more likely to stop trading after events because they are less resilient.

This is confirmed by an analysis of the vendors active during Operation Onymous. Users with lower ranks (higher reputations) were more likely to continue to trade.

It was also argued that losing reputation may indirectly diminish a user's capacity to trade by losing them trading partners. Some comments found in the subreddit discussions of the closure of *Alphabay* indicated that buyers were struggling to prove their reputation to vendors after losing an account.

Silk Road 2.0 data was used to evaluate Hypothesis 14 quantitatively. ERGM analysis was used to understand which qualities made vendors popular, where popularity was measured by the number and value of their transactions. One of the variables considered in this analysis was the reputation of the vendor. Whilst the reputation was shown to be a characteristic that influenced vendor popularity, looking specifically at the network in the weeks preceding the hack of *Silk Road 2.0* and immediately after this event (refer to Appendix H), increasing the reputation of a vendor does not increase the likelihood that someone will make a purchase from them. More importantly, the influence of vendor reputation remained unchanged during the event and for a month afterwards.

At this point, the influence of reputation on vendor popularity was no longer significant. This could potentially be because of the event and was a delayed response but could also be unrelated. Therefore, there is insufficient evidence to support Hypothesis 14 and further research is needed.

Events can cause users to lose product, either because the product is seized by a law enforcement or because a vendor does not send through an order. Figure 3.2 hypothesises that losing product can diminish a user's capacity to trade.

Insufficient data was available to evaluate if having parcels seized by law enforcement was sufficient to make users stop trading. However, analysis of the *Silk Road 2.0* dataset was used to understand how transactions informed buyers' decisions more generally.

A model was created to understand which elements of a transaction might cause them to stop trading. Whilst this showed that the higher a buyer rated a transaction the more likely they were to continue trading, there were an insufficient number of occasions in which the buyer did stop trading for the model to be considered reliable.

Further, it was shown that the proportion of a buyer's transactions that were negatively rated and their lifespan were not strongly negatively correlated.

There are lots of reasons a buyer might rate their experience negatively beyond having a parcel not arrive because it has been seized or the vendor they were trading with has been arrested. However, this data does imply that events that only result in buyers losing their product are unlikely to diminish their capacity to trade. Even though Hypothesis 15 was not directly tested, this analysis makes it seem less likely to be true.

When users perceive law enforcement to be behind events on the ecosystem, their fear of the repercussions from law enforcement may diminish their capacity to trade. To determine if this is true, Hypothesis 16 argued that proportionately more users would leave the ecosystem after Operation Onymous (an event conducted by law enforcement) than the Evolution Exit scam (an event which was not conducted by law enforcement). Further, this would only be the case for users who feel that the event may result in their arrest, i.e. their data has been collected in the event.

A direct comparison of Operation Onymous and the Evolution Exit Scam shows this to be the case. A greater proportion of affected users stop trading after Operation Onymous than the Evolution Exit Scam even though the proportion of unaffected users was similar. However, Operation Onymous did result in the closure of more DNMs and so, when considering this dimension, was a more impactful event and this could also be the reason why more users left.

Hypothesis 17 argued that if Operation Hyperion, which predominantly sought to remind users of the presence of law enforcement and did not make other impacts (e.g. did not remove trading platforms, accounts, coin, reputation, product, or trading partners for users), diminished a user's capacity to trade it would reduce the size of the ecosystem.

The effect of Operation Hyperion on the ecosystem population could not be measured quantitatively. However, the discussion of Operation Hyperion in the subreddits evaluated was much smaller than that of the closures of *Alphabay* and *Hansa*. Further, within that discussion the majority of contributors speculating on the consequences for affected users felt that there would be no repercussions beyond the letters/phone calls received and that these were felt to not be that harmful.

Therefore, the fear of law enforcement may have an ability to diminish the capacity of affected users to trade, but potentially only through particular kinds of interventions.

An indirect way in which the fear of law enforcement might diminish a user's capacity to trade is hypothesised to be through reducing trading partners. This is argued in Hypothesis 18.

After *Hansa* closed, users became concerned that vendor accounts on the marketplace Dream had been taken over by law enforcement. Though the number of accounts identified as being affected by this event was small, comments discussing it also discussed whether or not the whole site had been compromised and users were warned against trading with the compromised accounts because it would allow law enforcement to collect information about them. This shows that a fear of law enforcement has been discussed as a reason for not trading with users.

Finally, losing a trading partner in and of itself was hypothesised as diminishing a user's capacity to trade. Evidence to evaluate Hypothesis 19 was found by exploring buyer vendor relationships on *Silk Road 2.0*. This data shows that the whilst some characteristics of the vendor a buyer was trading with and how that buyer evaluated the vendor were a significant influence on the model, they were not as important variables as the buyer's own experience or the nature of their last transaction. This data implies that losing a trading partner does not actually fully diminish a buyer's capacity to trade.

To summarise the evaluation of these hypotheses, the user model figure 3.2 is reconstructed in figure 9.1. In this figure hypotheses that have been supported are shown with solid lines, hypotheses that are partially supported have dashed lines, and hypotheses that are not supported have dotted lines.

The research in this thesis supports the conclusion that losing an account and/or reputation can diminish a user's capacity to trade; that losing product and/or trading partners and/or being made more aware of the capacity of law enforcement may also diminish this capacity but further work is needed; and that losing coin and/or experiencing ecosystem instability does not diminish a user's capacity to trade.

Further, whilst a heightened fear or awareness of law enforcement may result in the loss of trading partners, losing an account or reputation does not necessitate losing trading partners. Finally, losing an account may result in the loss of coin and/or reputation but more evidence is needed to substantiate these claims.

9.1.2 Ecosystem Resilience

There were four hypotheses evaluating figure 3.1, the model describing how the events on an ecosystem affect its overall resilience.

Hypothesis 1 argued that the Evolution Exit Scam would have a bigger impact on the ecosystem than Operation Onymous because it removed more money from the ecosystem. However, the opposite was found to be the case as discussed in the previous section. This implies that the overall ecosystem is not necessarily impacted in a manner proportional to the amount an event impacts the cashflow.

Hypothesis 2 argues that heightened periods of arrests should diminish the capacity of the ecosystem more than heightened periods of parcel seizures. Precise data on the impact of neither heightened periods of arrests nor parcel seizures could be found.

However, within the discussion of Operation Hyperion, several contributors mentioned Operation Pangea, a pharmaceuticals based operation coordinated by Interpol (Interpol

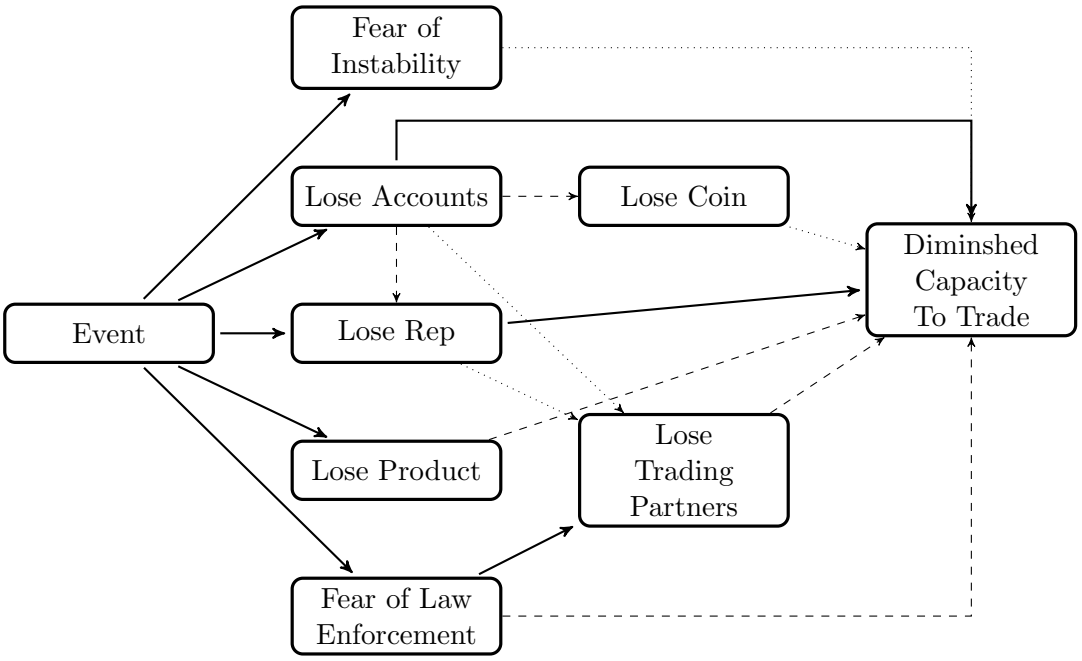

FIGURE 9.1: **Model for the Impacts of Events on Users** (solid lines represent supported hypotheses, dashed lines represent partially support hypotheses and dotted lines represent unsupported hypotheses)

(2018)). Operation Pangea involved several actions including arrests but was referred to in the forums exclusively in terms of seized packages, specifically those sent from Canada to the UK. This prompted several contributors to ask if Canadian vendors should temporarily stop trading with UK buyers.

Similarly, the actual arrests that occurred through Operation Hyperion were not discussed by contributors. Instead, the discussion of the mechanics of the operation was focused on the warnings received by users. For both operations, this is plausibly because the users who had been arrested were unable to contribute to the forums (though law enforcement announced they were making arrests in both operations).

This could indicate that contributors are actually more concerned about parcel seizures than arrests when those parcel seizures could affect them, but the arrests are happening to other users. Though this does not disprove Hypothesis 2 it indicates that more research is needed with more relevant data.

As Operation Onymous and Operation Bayonet removed trading platforms they directly reduced the trade occurring in the ecosystem where Operation Hyperion did not. Therefore, to evaluate if a reduction in trade diminished the capacity of the ecosystem, Hypothesis 3 argued that these events would have a bigger impact than Operation Hyperion.

Operation Onymous and Operation Hyperion could not be directly compared with the data collected. Operation Bayonet and the closure of *Hansa* were qualitatively compared to Operation Hyperion. In the former operations significantly more comments and a wider variety of impacts were recorded. Whilst the impacts of Operation Hyperion and the closures of *Alphabay* and *Hansa* are not directly compared, the impacts discussed relative to the latter are described more negatively. This implies that, from a qualitative perspective, Operation Hyperion was not considered as impactful as Operation Bayonet and the closure of *Hansa*.

This perspective is supported by comments about Operation Hyperion arguing that this operation was not as severe as other operations that proceeded it.

Finally, to understand if reducing the perceived convenience of the ecosystem diminishes its capacity, the comments after *Alphabay* and *Hansa* closed were read to understand how contributors perceived the convenience of the ecosystem during these events. Multiple comments expressed the sentiment that the ecosystem could not provide a previously experienced variety of products, that there were a heightened amount of scams, and users could not trust surviving trading platforms and that, therefore, it was not possible to trade on the ecosystem or that it was recommended to wait for things to return to normal before engaging in purchases. It is therefore argued that the qualitative analysis of Operation Bayonet and the closure of *Hansa* strongly support Hypothesis 4, that a perceived reduction in convenience is linked to a perceived diminished capacity.

The model of ecosystem resilience is redrawn as figure 9.2. The relationships that were confirmed by the research presented in this study are left as full lines, those that are partially supported are redrawn as dashed lines, and those that are not supported are drawn as dotted lines. Some are highlighted red to make multidirectional relationships more clear. The conclusions outlined in figure 9.1 have also been incorporated into the model

This model illustrates that removing trading platforms and disrupting the convenience of the ecosystem do diminish its capacity, however simply removing coin or from the ecosystem or reducing its population does not appear to diminish its capacity. Therefore, it is argued that the ecosystem is resilient to law enforcement operations that fail to remove trading platforms or disrupt the convenience of the trade, for example Operation Hyperion.

A summary of the hypotheses and whether or not they have been supported is given in table 9.1.

TABLE 9.1: A Summary of the Conclusions on Each Hypothesis

Hypothesis	Description	Supported
1	The greater the amount of money lost in an event, the bigger the event has.	Unsupported
2	Events that reduce the size of the population have a bigger impact than those that do not.	Unsupported
3	Events that reduce the amount of trade have a bigger impact than those that do not.	Supported
4	If an event reduces the convenience of trade it has an impact.	Supported
5	The more DNMs that a user has witnessed close, the more likely they are to leave the ecosystem after an event.	Unsupported
6	Losing an account makes users stop trading.	Supported
7	Users with more accounts are less likely to stop trading after losing an account.	Supported
8	Users are unable to retain their reputation after losing an account.	Partially Supported
9	Users are unable to maintain relationships with trading partners after losing an account.	Unsupported
10	Users who lose accounts also lose money.	Partially Supported
11	If a user loses money in an event they are more likely to stop trading.	Unsupported
12	Users are more likely to stop trading if they lose an account and money than if they just lose money.	Supported
13	Users with lower reputations are more likely to stop trading after an event.	Supported
14	Users who lose reputation are unable to keep their trading partners.	Unsupported
15	Users who lose product in an event are more likely to stop trading.	Partially Supported
16	Events that involve law enforcement are more likely to make vendors stop trading than events that don't.	Partially Supported
17	Events designed to deter users by increasing the awareness of law enforcement result in users stopping trading.	Partially Supported
18	Events designed to deter users result in users losing trading partners.	Supported
19	Users who lose trading partners are more likely stop trading.	Partially Supported

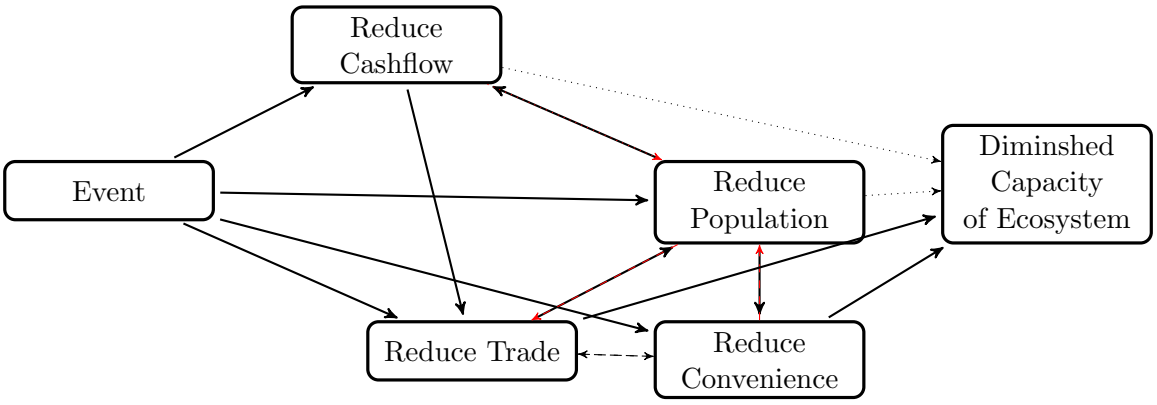

FIGURE 9.2: **Hypothesis Model for the Impacts of Events on Ecosystem** (solid lines represent supported hypotheses, dashed lines represent partially support hypotheses and dotted lines represent unsupported hypotheses, red has been used to highlight two hypotheses discussing the same relationship)

9.2 Limitations

For each of the studies, decisions had to be made in the data collection and preparation methods which may have implications for the results and their interpretations. These limitations and how they affect the interpretation of the results are presented for each study.

9.2.1 Cross Market Study

The data used in this study is incomplete (Branwen et al. (2015)). For each individual DNM included in the study, vendor, listing and especially review pages are missing and, therefore, the size of each DNM is underestimated. Further, the dataset did not capture every DNM operational during the observation period and not every DNM included in the dataset could be used.

This affects the results in several ways. Firstly, if the DNMs omitted from the dataset were popular marketplaces that vendors migrated to after Operation Onymous or the Evolution Exit scam then the measurements taken of the population either side of these events would over estimate their impact. This is unlikely to be the case, however, as it is believed that the major operational DNMs are present in the dataset.

Secondly, the number of accounts controlled by each vendor may be under reported as not all of their accounts are captured in the dataset. Similarly, some vendor lifetimes may be underestimated. However, as this is the case for all vendors, conclusions made about how these attributes affected whether or not vendors continued to trade after events may still be approximated.

Finally, the review data was too impaired to be used for any of the DNMs. As a result, the method of approximating the number of sales as employed by Christin (2013) and Soska and Christin (2015) could not be used. The size of the vendor population was therefore the only available measure of the impact of events. This does not give a complete picture as some events may not reduce the population but still reduce the amount of purchases taking place.

As the public data only provides information on vendor profiles, buyer responses to events cannot be measured. The analysis of the Silk Road 2.0 dataset shows that buyers responded to events in much greater proportions than vendors. Potentially, therefore, the results of the Cross Market Study are underestimates and the events had bigger impacts on the population.

It could also be the case that buyers are more likely to leave the ecosystem after an exit scam than a law enforcement operation which would change the conclusions of the Cross Market Study. However, this could not be evaluated with either dataset.

Capturing the true size of the ecosystem requires knowing which vendor accounts are run by the same person or group of people. Whilst the techniques used within this study link more accounts than other, comparable studies, the lack of ground truth makes evaluating them difficult.

If vendor account links have been missed then the size of the ecosystem has been over reported. This also means that vendors who did continue to trade but with an account that could not be connected to them were counted as not continuing to trade. This overestimates the impact of the events evaluated.

Alternatively, if vendor accounts are incorrectly linked because the methodology employed cannot distinguish between a true account and an impersonation, for example, then the ecosystem population has been underestimated. Also given that, after large DNM closures, some users steal the usernames of others to hijack their reputation it is also possible that some vendors were recorded as continuing to trade after the events measured when they, in fact, did not.

As the same methodology was employed for the extent of the measurement period, these errors could have been applied consistently. Therefore, even if the precise impacts of Operation Onymous and the Evolution Exit Scam cannot be measured, they can still be compared relatively.

An additional methodological decision was the way the reputation measures were compared. This approach, outlined in Section. 5.1.4 converts the specific reputation measure of each DNM into a ranking designed to encapsulate both how popular the vendor was

and also how reputation was distributed across the marketplace. Vendors with the highest reputations were ranked in the top places but this meant less if many vendors on the DNM received similarly high reputation scores.

Whilst this method allowed for the reputation of vendors operating on different DNMs to be compared, it did not take into account the overall size of each DNM. Further, as different DNMs were valued differently, a high reputation on one might not be equivalent to a high reputation on another. Finally, the reputation scores may not be an accurate measurement for how the vendor is perceived.

As a result, the analysis on the influence of vendor reputation on popularity may produce different results if a different methodology is employed.

Finally, as this study is purely quantitative, the results can only express measurable changes in the population. The explanations given for why these changes occurred are informed by existing research but cannot be shown to be definitively correct. For example, when examining the reduction in population after the Evolution Exit Scam, this study cannot show that the accounts that closed were closed because of the exit scam.

9.2.2 Silk Road 2.0 Study

The dataset of vendor and buyer activity on the DNM Silk Road 2.0 is limited to users considered to be in the UK. Therefore, conclusions made with this dataset may not be transferable to other locations.

In addition, the dataset does not provide a complete picture of the activity. For non-UK users being traded with, their profile attributes (such as length of lifetime, total number of sales, etc.) had to be estimated and were likely underestimated. As much as possible, analysis was limited to UK users only in order to minimise the use of this estimated information. However, even when analysing UK users, attributes from their trading partners had to be evaluated. Therefore the ERGM analysis used to understand what makes vendors popular, will have preferenced UK based vendors and can only be used as an approximation.

Users may have had more than one account on the site but analysis to link accounts was not conducted. This means that the figures quoted for the size of the marketplace are potentially an overestimate. If users responded to the events evaluated, such as the hack or a vendor they were trading with leaving the market, by creating a new account and closing their existing one then this also has implications on the impact of these events. However, it is unclear why that behaviour would be likely.

As with the Cross Market Study, the analysis of Silk Road 2.0 was purely quantitative. When investigating what makes vendors popular or caused users to close their accounts it could not be shown why specific users enacted this behaviour, only that it occurred.

9.2.3 Reddit Forums Study

With regards to the data collected from Reddit, there are both specific limitations and more general limitations associated with working with forum data.

The specific dataset is limited because many of the actual posts are missing and therefore context had to be determined from the comments alone. There is the possibility that comments were mislabelled because a similar sounding context was actually being discussed. For example, if a comment expressing an opinion about a site being down was presumed to be about *Alphabay* but was actually about the Dream marketplace. To minimise this, as much as was possible, comments were only labelled if context could be confidently determined.

The way that the Reddit data was collected also meant that not all comments posted on the same thread were given the same thread ID and therefore couldn't be linked. This primarily affected the statistics measuring the network, such as the density and transitivity. These measures are underestimates as a result.

A large proportion of users in both DNM subreddits either deleted their accounts or had them deleted. This significantly reduced the amount of available information on contributors. Most contributors only made one post or comment but it is highly likely that forum users changed their accounts and so their full activity could not be assessed.

As there were multiple Reddit forums on similar topics and only two were qualitatively analysed, it cannot be assumed that the full extents of Operation Hyperion and Operation Bayonet were captured in the study. Similarly, the qualitative impact evaluations were conducted on a subset of the forums produced through keyword searches. Though the aim was to gather more comments, even if not all were relevant, it is possible that relevant comments have not been read.

More generally, when working with forum data, analysis can only assess what is being said by contributors, and this is not necessarily an accurate representation of what has happened. It is possible that contributors exaggerated or falsified their claims and that many DNM users did not use the forums to discuss their experience of the events. As a result, conclusions from this study are limited to what contributors of the forum report has happened, rather than what has happened.

9.3 New Findings and Future Work

In addition a number of new findings were identified that were not predicted in Chapter 3. These new findings, and the potential areas for future work that they inspire, are discussed in the following section.

A key discussion that occurred in the Reddit forums was how *Alphabay* and *Hansa* were closed. Some contributors believed official law enforcement statements that explained how they exploited human errors to identify the administrators. Others, however, were concerned about a Tor exploit or similar structural vulnerability. These contributors also seemed more worried about further attacks and site closures.

As Operation Onymous involved closing 414 `.onion` sites it was more widely believed to be the conducted by exploiting a Tor vulnerability. A qualitative exploration of relevant forums in which Operation Onymous is discussed could identify if contributors similarly felt operations which employ such a tactic are more concerning than when human error is exploited. Additionally, a quantitative assessment of the impact of Operation Bayonet and the closure of *Hansa*, at least on the vendor population, would enable a direct comparison between these closures and Operation Onymous. This could be used to compare which type of operation is more effective.

The analysis of Operation Onymous shows that vendors mostly migrate to one marketplace when the DNMs they are trading on are shut. This fact was used effectively when *Alphabay* and *Hansa* were closed in 2017. However, this operation complicated the discussion about where to move to next after *Hansa's* closure. Contributors became concerned that, because *Hansa* had been recommended as the site to move to post-*Alphabay*, the new most highly recommended sites were also honeypots.

This observation implies that law enforcement operations can spread distrust within the forums as well as having an impact on DNMs. Quantitative analysis of the ecosystem population after the closure of *Hansa* could show where vendors (and buyers, if data is available) moved and if they did so seemingly at random or, at least, less consistently than after Operations Onymous and Bayonet.

When *Hansa* banned fentanyl the action was welcomed by some contributors and rejected by others. The contributors who approved of the ban both argued that the drug was too dangerous and that it was attracting too much attention from law enforcement, i.e. the arguments were both principled and pragmatic. Despite those in the community who disapproved, the action has been replicated on other DNMs.

The ban on fentanyl potentially shows that the ecosystem is prepared to compromise on activities they are concerned are high priority for law enforcement in order to better

protect themselves. This could be tested through a controlled experiment to target a different product.

Fentanyl was not the only product discussed as attracting too much attention, others included child pornography (which is widely banned on DNMs), opioids more generally, and fraud related items. If, for example, fraud became a widely publicised priority DNMs could be monitored to see if they either discuss or actually ban products associated with fraud.

Alternatively, a geographical analysis of DNM vendors and their products compared to local law enforcement objectives could measure if users are less willing to trade in products that they know are being highly investigated in their area.

This thesis has identified several new findings that potentially explain user behaviour and reveal new insights into the impact of law enforcement operations. In order to more fully understand them, a quantitative evaluation of Operation Bayonet and a qualitative exploration of relevant forums in the wake of Operation Onymous are recommended. As is an investigation into the types of products users are willing to trade in and why.

9.4 Recommendations for Law Enforcement

An assessment of three law enforcement operations (Operation Onymous, Operation Hyperion, and Operation Bayonet/the closure of *Hansa*) gives insight onto the different impacts that such operations can have on the ecosystem. Operation Onymous and Operation Bayonet/the closure of *Hansa* had an immediate impact on the ecosystem and on users because they resulted in accounts being closed.

When users lost access to their accounts, many lost money, which was reported to be in the range of five figures for some users. Additional reported losses were of reputation which could help establishing trust on new markets and of trading partners who could not be contacted either for the purpose of following up on a transaction or to establish a new trade. Analysis on the interactions between vendors and buyers on *Silk Road 2.0* shows that loss of a vendor can be enough to draw buyers away from trading on a DNM.

A byproduct of this type of operation is the turmoil created by mass migration of displaced users. When many users who have just lost an account want to continue trading but do not have established accounts to switch to, such as after the closure of *Alphabay*, the population is observed to overwhelm surviving marketplaces who are unprepared to process the influx of new users. Within the forum discussions post Operation Bayonet, contributors warned others of users exploiting this chaos for example by creating phishing

links to fake websites and impersonating *Alphabay* vendors with the view of scamming their customers.

It should be noted that this turmoil appears to have been short lived with users eventually settling on Dream Market. A quantitative analysis of the operations over a longer observation period could more accurately assess the impact of mass user migration.

In addition to closing accounts, these law enforcement operations can also lead to the collection of data on users. Forum discussions, especially after contributors learned of how *Hansa* been used as a honeypot, show concern for how this data might be used to investigate and arrest users. The immediate response to this impact was for contributors to recommend users "lay low", at least for a short period. This element of an operation changes the impacts of an operation, if not making it more impactful than a comparably sized exit scam. This can be seen in the discussion of *Alphabay* before and after the FBI announced their operation as contributors switch focus from financial loss to if their data has been collected and what this means. It can also be seen that Operation Onymous has a proportionally bigger impact on the affected population than the Evolution Exit Scam.

There is also a less immediate potential impact when law enforcement have processed the information collected in an operation either in the form of arrests or warnings as in Operation Hyperion. Arrests obviously have consequences for the users directly involved, but can also result in consequences for other users. If the person arrested is a vendor their buyers lose a trading partner and could also lose a product they have purchased that hasn't been shipped out. If the vendor does not properly secure information on their buyers then they also place their buyers at risk of being investigated and arrested.

The forum discussions surrounding Operation Hyperion imply that other actions taken by law enforcement that do not result in arrests, e.g. warnings, are easily dismissed by users. This is especially the case when the information that law enforcement reveal they are operating on is several years old. An assessment of this data and the discussions following the closures of *Alphabay* and *Hansa* give the impression that contributors are concerned about immediate arrests, as opposed to those in the distant future.

These law enforcement operations did not affect all users equally. Analysis on the type of users who stopped trading vs those who continued to trade after Operation Onymous shows that users who are directly effected are much more likely to stop trading. In addition, the vendors most likely to cease trading were those with lower reputations i.e. those who are less popular.

From the discussions on Reddit after the closures of *Alphabay* and *Hansa*, contributors expected ensuing investigations to be directed towards vendors and buyers involved in large scale purchases (implying that they were buying wholesale and reselling).

Additionally, contributors felt that users were able to protect themselves from the repercussions of law enforcement operations by employing good OPSEC and being cautious[1]. For example, if you do not store coin on a DNM, you can't lose it when the site closes and if you encrypt all of your communications and do not rely on the site's autoencryption system, law enforcement won't find information about you even if they seize the site servers.

This was part of a wider narrative about how users could protect themselves from law enforcement operations. This narrative included discussions about how law enforcement were able to close *Alphabay* and *Hansa*. The official explanations given by the FBI and Dutch Police, respectively, relied on exploiting admin errors. Contributors who thought that law enforcement actually exploited the Tor network or used a different technological approach were concerned that more DNMs would be closed as part of a continuing operation.

This assessment of the three law enforcement operations studied in this thesis informs the following set of recommendations for law enforcement:

- the objective of an operation should be to actively make it harder to trade (e.g. by making an arrest, closing an account, or shutting down a platform) rather than passively making it harder to trade (e.g. by warning users or otherwise trying to deter them);
- the actions which actively prevent users from trading need to take place immediately (as opposed to in several years);
- either operations should be designed to affect users as universally as possible (e.g. when making arrests do not only target large scale vendors but all users at all activity levels) or should specifically target the most desirable users (this may mean not employing the current model of targeting whole DNMs as these large scale operations appear to have greater impacts on less important users);
- where possible operations should expose security flaws in the system and, at least outwardly, not appear to rely on human errors;
- operations should be regular, both in terms of occurrence and impact, so that the presence of law enforcement is consistently felt, Reddit contributors already expect

[1]The concept of Restrictive Deterrence was raised in the viva to describe this phenomenon. Restrictive Deterrence models how offenders don't necessarily stop offending after a law enforcement intervention but, instead, change their behaviour to avoid being (re)caught. There is evidence to suggest that this change in behaviour makes some offenders more vulnerable to arrest because they have adopted techniques that they are less practised at (Paternoster (1989)). Therefore, more research is needed to see if employing more OPSEC actually afforded users greater protection.

an operation annually and this is felt to be sufficiently intermittent that the impact can be absorbed however they still need to remain unpredictable - Operation Bayonet and the closure of *Hansa* demonstrates how operations in several stages can draw out and amplify an impact;

- when closing DNMs in an operation, the more sites and the bigger the sites closed, the better;
- even though users seem to think law enforcement operations are fundamentally illegitimate, progress is made when law enforcement goals align with the general ecosystem opinion, this can be seen with the ban of fentanyl which happened on *Hansa* and was replicated on other DNMs.

Chapter 10

Conclusion

This thesis has presented quantitative and qualitative analysis of three datasets related to DNMs active between 2013 and 2017. It has used this analysis to explore the research question *Is the DNM ecosystem resilient to law enforcement interventions?* and considered this question from a user and ecosystem perspective. This chapter will present a summary of the work and conclude on this research question.

The resilience of the ecosystem was defined as a measure for how able the ecosystem facilitates trade despite events such as DNM closures, hacks, and exit scams. In addition, the resilience of ecosystem users was defined as their ability to continue to trade successfully whilst experiencing the same adverse events.

In Chapter 3, two models were presented that articulate how an event might diminish the capacity of a user to trade or the capacity of the ecosystem overall. In addition, 19 hypotheses were given to test these models.

Existing literature has explored the impact of individual events on the ecosystem and its population, however, this thesis is the first body of work which has tried to model common factors that determine the ecosystem's response to such events. Further, this thesis asks a new research question and presents new work by strategically comparing the outcome of different events in order to understand the resilience factors of the ecosystem and its population.

To answer the research question, three datasets were collected. These are described in Chapter 4. The studies conducted on each dataset are described in Chapters 5,6,7 and 8. These studies are briefly summarised below.

10.1 Cross Market Study

The first dataset used is a large, public dataset containing scrapes of 87 DNMs collected between July 2013 and July 2015. This dataset was collected by the independent researcher Gwern Branwen and has been used extensively within this field (Branwen et al. (2015)). Despite this, this thesis is the first piece of research to thoroughly evaluate the completeness of this dataset.

A novel technique was presented which involved comparing the pattern of growth in the observed dataset and the expected pattern of growth formulated from additional information found within dataset and external datasets. This method resulted in the conclusion that 52 datasets were not complete enough for analysis. Additionally, all of the review data for vendors and products was either considered too incomplete or unverifiable.

This dataset was demonstrated to be an underestimate of the ecosystem both in terms of the number of DNMs represented and the amount of information it contained on each DNM. However, it was also shown that the underestimate is likely to behave in a similar manner to the actual ecosystem over time, in terms of the number of DNMs, vendors, and products.

During the observation period for this dataset, the law enforcement intervention Operation Onymous took place. As did the Evolution Exit Scam, these two large events were key examples outlined in Chapter 3 as events that could impact on the resilience of the ecosystem and its users. As such, this data was collected to understand how the ecosystem population responded to these events.

To do this some new methodological processes were developed. The first pertained to matching vendor accounts on different marketplaces. In this study existing methods of matching vendor accounts by comparing their usernames and PGP keys were improved upon by comparing profile descriptions using TF-IDF analysis and comparing products sold. These new approaches were shown to identify vendor account pairings that were not captured by reconstructions of approaches taken in other work.

The second was an approach for comparing the reputation of vendors operating on different DNMs. There were multiple metrics of vendor reputation contained within the dataset and, as such, a method was created to compare these different values in order to understand the reputation of a vendor across the observed ecosystem.

Vendor reputation values were translated into rankings which captured how their reputation compared to that of every other vendor on the DNM relative to how many vendors were active on it. Vendors with higher reputations had lower rankings. This

approach allowed for two vendors operating on different DNMs to be compared by their reputation even if those DNMs had different metrics for reputation.

It was found that Operation Onymous predominantly affected the vendors who were only trading on DNMs closed in the operation. The next most affected group was vendors who were trading on at least on DNM that was closed and one that was not closed. The population not trading on any DNMs closed in Operation Onymous were, by contrast, barely affected at all.

Further, the types of vendors most likely to cease trading were vendors with higher reputations and vendors who had witnessed multiple DNM closures during their lifetime. Interestingly, vendors were also more likely to continue trading if they were younger and had had less accounts. These results were surprising as it was expected that users with more accounts would be more able to continue trading after the operation.

An evaluation of the Evolution Exit Scam was conducted and the results were compared to Operation Onymous. It was found that, though the proportion of unaffected users who continued to trade was similar after both events, a larger proportion of affected users continued to trade after the Evolution Exit Scam than Operation Onymous. This was the case despite the fact that dramatically more coin was lost in total and as an average per vendor after the Evolution Exit Scam.

The users who continued to trade after the Evolution Exit Scam had fewer accounts, higher reputations, more PGP keys and had experienced more DNM closures, i.e. were similar to those who continued to trade after Operation Onymous. It was also found that vendors who lost an account in Operation Onymous who were active during the Evolution Exit Scam were more likely to stop trading.

Finally, both the Evolution Exit Scam and Operation Onymous were compared to the closure of the DNM *Darkbay*. *Darkbay* closed amicably and provided an opportunity for users of the site to conclude their business and withdraw coin before it did so. Because of this, and the fact that *Darkbay* was much smaller than *Evolution* and the combined size of the DNMs closed in Operation Onymous, it was anticipated that it would have a smaller impact on the ecosystem. This was found to be the case.

However, though *Darkbay* had a much smaller impact overall, and despite the fact that the proportion of vendors not trading on *Darkbay* who stopped trading was similar in size to the equivalent populations after Operation Onymous and the Evolution Exit Scam, a smaller proportion of affected vendors continued to trade after the DNM closed than for either of the other events.

10.2 Silk Road 2.0 Study

The second dataset was collected from the servers of the DNM *Silk Road 2.0* taken during the observation period 4 November 2013 - 6 November 2014. This dataset contains profile and transaction data on 24,393 accounts identified as belonging to UK users as well as transaction data for an additional 17,800 non UK users. The dataset capture a total of 190,802 transactions.

The dataset was presumed to be complete for UK users, but was limited for all other user locations. Despite this, it is the most complete dataset containing buyer data that has been utilised within the field. The data was received in the form of .xml files and so no additional preparation was needed except to anonymise the data by replacing usernames with ID numbers.

The dataset was compiled into a network with users represented as nodes and transactions as directed edges. The data was predominantly analysed using Network Analysis techniques such as the density, reciprocity, transitivity, and centralization of the network and ERGM analysis. The methodology was modelled on research by Duxbury and Haynie (2017).

The measure of density was considered to be inappropriate for this network as it implies that any user (be they buyer or vendor) may buy from any other user. This is not the case for this dataset as most users either only buy or only sell. Instead a new measure of density was proposed which measures the number of edges in the network as a proportion of the number of possible buyer/vendor pairings.

Additionally, a method of assigning locations to users using the information captured in their postage data was developed. This technique builds upon existing work which uses the shipping information published on vendor profiles to assign them locations.

Analysis of the data found that the network modelled on *Silk Road 2.0* had a low density, transitivity and reciprocity implying that the users are not greatly connected and that connections are generally one way. Further, connections were clustered around specific vendors, rather than buyers.

The size of the network was examined over its lifetime. This showed one dramatic change in population size, concluded to have been caused by the hack of the site which took place in February 2014. No other significant events were identified, however, the population fluctuated considerably at the beginning and end of the DNM's lifetime potentially obscuring the effect of other events. It was also found that events had a greater impact on the amount of trade, rather than number of accounts.

To understand the relationships between buyers and vendors, ERGM analysis was used to evaluate different variables which might make vendors popular. Some variables (the vendor age, reputation and diversity of products sold) had a marginal positive impact on the probability of a tie being formed for most snapshots whereas others mostly did not affect the probability of a tie being formed. This differed from the results found by Duxbury and Haynie (2017) who found that vendor reputation, in particular, had the ability to significantly increase chances of a tie being formed.

Finally, it was asked whether or not vendors leaving the DNM would prompt buyers to also leave. This was explored because, if buyers did leave when a vendor they were trading with left, it could imply that buyers preferred trading with particular vendors and were either prepared to follow them to different trading platforms or did not want to trade at all if they were no longer able to trade with them on *Silk Road 2.0*.

Logistic Regression Analysis of incidences in which vendors closed their accounts and then their buyers either stopped trading or continued trading showed that the characteristics of the specific vendor a buyer was trading with, and their opinion of that vendor, were far less influential variables than the buyer's experience and the nature of their last transaction.

Logistic Regression Analysis was also used to understand after what kinds of transactions buyers were more likely to leave the site. This showed that more experienced buyers were more likely to continue trading after a transaction, regardless of its nature, but more importantly rating the transaction positively and a history of trading with that particular vendor before were the most important determinants of the buyer making more purchases.

Though both models built were accurate in terms of their ability to predict when buyers would stop or continue to trade, there was a high skew towards buyers continuing to trade in both which likely improved the models ability to guess regardless of the variables considered.

10.3 Reddit Forum Study

The third and final dataset contains posts and comments from three Reddit forums. These are the subreddits */r/darknetmarkets*, */r/dnmuk* and */r/Ebay*. They were collected from the public repository created by Stuck_In_the_Matrix (2015) via Google BigQuery. Though */r/darknetmarkets* and */r/dnmuk* have now been closed by Reddit administrators, they were both active during the observation period of 1 September 2016 to 30 November 2017.

Posts and comments for each dataset were collected separately and then aggregated by matching the post ID numbers. The datasets were evaluated for completeness by counting the number of deleted or removed comments, posts and users and by counting the number of comments that could not be attributed to a post.

/r/darknetmarkets contained approximately 324,120 posts and 572,585 comments with 97% of the comment threads attributed to posts not found in the dataset. This was either because the posts were not collected or the comments were direct replies to other comments and could not be matched to them as this information was not contained in the dataset. 4,088 (38%) of the posts found in the dataset had been deleted and the recorded number of comments was predicted to be 39% of the total subreddit. This subreddit had 417,701 contributors, 69,858 of which had deleted usernames.

The subreddit */r/dnmuk* contained 168,873 posts and 281,248 comments. 99% of the posts were artificially created from comments and, of an expected 21,997 comments, only 39% were found. 26% of the posts found had been deleted. There were 200,792 contributors found in the dataset, of which 27,643 had had their username deleted.

The subreddit */r/Ebay* contained 54,606 posts and 87,264 comments. 16% of the posts had been deleted and 91% were created from comments. Of an expected 18,770 comments 46% were found. There were an estimated 64,428 contributors in the dataset, 5,798 of which had had their username deleted.

As in the *Silk Road 2.0* study, a network was created such that nodes represented contributors and edges represented when one contributor commented on the post of another. Quantitative analysis on each forum showed that the networks created were sparse and that contributors rarely made multiple contributions using the same account. It was also shown that a higher proportion of comments and users were deleted in the Dark Web forums than on */r/EBay*.

The Dark Web forums were used to explore Operations Hyperion and Bayonet and the closure of the DNM *Hansa* through both quantitative and qualitative approaches.

The impact of each event was analysed quantitatively by measuring the change in size of the DNM related forums over the weeks in which they took place. This analysis shows a large increase in activity during Operation Bayonet and after the closure of Hansa but no such increase is visible during Operation Hyperion. Interestingly, not only do the number of comments and posts increase during Operation Bayonet but so does the number of contributors, implying that more users sought out the forums for information.

Contributor responses were then explored qualitatively by reading relevant posts and comments to identify themes and topics using a combination of Grounded Theory and Directed Content Analysis.

Analysis of Operation Hyperion related posts and comments reveals that contributors were not concerned about Operation Hyperion. It was compared to previous operations and considered to be of a lower impact with contributors arguing that those targeted would not face any real repercussions. Users who received letters or phone calls were blamed for having poor OPSEC which led to their identification.

In contrast, the consequences of Operation Bayonet and the closure of *Hansa* were discussed in detail. Contributors listed financial losses, some of which were described as extremely high, as well a great concern for how law enforcement would use any data seized in the operations. There were more discussions of financial losses after the closure of *Alphabay* and more discussions about the repercussions of data seizure after the closure of *Hansa*.

As with the discussion surrounding Operation Hyperion, contributors felt that those who were most at risk or who had lost the most were themselves to blame for having poor OPSEC and not taking the necessary precautions when trading. Many more comments during these operations warned other users to stay away from the ecosystem than after Operation Hyperion.

Many contributors were concerned about how law enforcement had conducted each operation and a large theme on this topic was that official stories were deliberate misinformation spread to hide the real tactics used. Where contributors discussing Operation Hyperion felt that law enforcement were incompetent or failing to tackle to ecosystem, some contributors were impressed by the combined operations against *Alphabay* and *Hansa* and were concerned for the future of the ecosystem. Despite this there were still many contributors who argued that the operations had not been successful and that the ecosystem would recover.

10.4 Findings

These studies were used to answer 19 hypotheses that predicted how different ecosystem events might impact the ecosystem and its population. The events evaluated were DNM Shut Downs, User Warnings, Parcel Seizures, Closures, Exit Scams, and Hacks. These hypotheses described how reducing the cashflow, trade, population, and convenience of the ecosystem might diminish its capacity to trade, thus making it non resilient and how if an individual user lost an account, reputation, product, coin, and/or trading partners

or developed a fear of instability or fear of law enforcement they may also become less resilient.

This thesis argues that events that result in a fear of instability or the loss of coin do not appear to diminish a user's capacity to trade. However, events that reduce users' reputation or result in them losing an account can diminish their capacity to trade.

More research is needed to understand fully the impact of events that cause users to lose product or trading partners as the claim that these events reduce a user's capacity to trade was only partially substantiated.

Finally, this thesis found evidence in support of the argument that a heightened fear of law enforcement diminished a user's capacity to trade. However, evidence was also found demonstrating that this is not always the case. Therefore, users are not always resilient to events which increase the fear of law enforcement.

Based on the specified impacts of each event outlined in Chapter 3, this leads to the conclusion that the population is resilient to User Warnings and partially resilient to the remaining events.

When looking at the ecosystem, it was found that events that reduced the population did not actually diminish the capacity of the ecosystem. Therefore, even though individuals may not be resilient to the events evaluated in this thesis, the ecosystem as a whole has been.

The ecosystem was also found to be resilient to events such as Hacks that reduced the cashflow of the ecosystem without other impacts. However, when events (such as DNM Shut Downs, Closures, and Exit Scams) reduced trade, for example by removing trading platforms, or reduced convenience, for example by increasing the amount of scams taking place, then the capacity of the ecosystem was diminished. It is therefore argued that the ecosystem is less resilient to these types of events than others.

Each of the datasets employed in this research is incomplete. This means that the full extent of the impacts of each event evaluated may not be fully seen. In the case of the Cross Market Study, the part of the ecosystem not captured in the dataset could reveal a more resilient user population or it could show that sales or the buyer population were effected more considerably than the number of vendor accounts. For the analysis of the Reddit forums, posts and comments omitted from the dataset may document more impacts of Operations Hyperion and Bayonet or a greater proportion of contributors stating that they had not been affected.

Further all of the events could only be analysed either qualitatively or quantitatively. Therefore, when measuring the number of vendors who left after Operation Onymous,

for example, or the number of buyers who stopped trading after a vendor they had purchased from left, it cannot also be shown if these users stopped trading for the reason given. Alternatively, the qualitative evaluations of Operation Hyperion and Operation Bayonet and the closure of *Hansa* cannot give the scale of of these impacts.

As such, this thesis is unable to make conclusions on which events the ecosystem is resilient to. Instead it compares different approaches and can more confidently make conclusions on which events the ecosystem is less resilient to. For example, the ecosystem was more resilient against Operation Hyperion than Operation Bayonet and the closure of *Hansa*. It was also more resilient against the Evolution Exit Scam than Operation Onymous.

Despite its limitations, this work has made positive contributions to both the academic literature on DNMs and current law enforcement approaches to combating them. The main contributions of this work are the new methodology used to evaluate the Gwern Branwen dataset, the acquisition of the *Silk Road 2.0* dataset, and the assessment of the operations Hyperion and Bayonet. Each of these contributions have implications for existing work, as they can be used to better contextualise or more critically understand the reliability of research using the same or different datasets or other methodologies.

Further, as much as possible, the results of this thesis have been discussed in relation to current law enforcement approaches to combating DNMs. The DNM ecosystem continues to grow and has facilitated the sale of priority drugs such as fentanyl. As such, this thesis contributes to a growing body of work that can help to direct law enforcement approaches in a more efficient and effective manner. By directly comparing different approaches, this work highlights how targeting specific types of actors on the ecosystem can be more effective than other approaches and that ecosystem users are not necessarily deterred unless they are impacted by an operation. These findings can help law enforcement when designing future operations.

Appendix A

Validation Metrics

1776

The number of products advertised on the site compared to the number of unique products collected.

Abraxas

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Agora

Lower estimates for number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Alpaca

The number of products advertised on the site, measured first as the stated number of available products and then as the sum of categories, compared to the number of unique products collected.

Alphabay

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

AmazonDark

The number of products advertised on the site compared to the number of unique products collected.

Anarchia

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Andromeda

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Area51

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Black Services Market

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Bloomsfield

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Bluesky

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Breaking Bad

None

Cannabis Road 2

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Cannabis Road 3

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.
textbfCantina

The number of products advertised on the site compared to the number of unique products collected.

Cloud 9

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values. The number of vendors advertised on the site compared to the number of unique vendors collected. The number vendors advertised collected from “ship from” data.

Cryptomarket

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Dark Bay

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Dark List

None

Dark Net Heroes

None

Deepzon

None

Diabolus

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Dogeroad

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Dreammarket

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Drugslist

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

East India Company

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Evolution

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Freebay

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised taken to be the value assigned to "All Items".

Free Market

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised taken to be the value assigned to "Total Listing Count".

Grey Road

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Haven

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Horizon

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Hydra

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Ironclad

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Kiss

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Middle Earth

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised as stated on index page.

Nucleus

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Onion Shop

None

Outlaw Market

The number of vendors compared to the number that are in the drop down menu on the vendor page.

Oxygen

None

Panacea

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values. The number of vendors compared to the stated number of vendors collected from the index page describing the shipping information.

Pandora

None

Pigeon

None

Pirate Market

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Poseidon

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Silk Road 2

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Silk Road Reloaded

None

Silk Street

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

The Majestic Garden

None

The Market Place

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised taken as number of drugs advertised plus number of technology products advertised. Category information collected for each listing.

The Real Deal

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Tochka

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Tom

None

Topix 2

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised the stated number of drugs.

Torbay

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Tor Bazaar

None.

Tor Escrow

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values

and category information collected for each listing.

Tor Market

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

Tortuga 2

None

Underground Market

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values.

White Rabbit

The number of products advertised on the site compared to the number of unique products collected. The number of products advertised calculated by adding category values and category information collected for each listing.

Appendix B

Username Matches

“reepee911”, “evopee911”, “palpee911”, ”aaapee911”, “diapee911”, “raxpee911”, “aaapee911”, “keypee911”, “diapee911 premium”

“FriendlyNeighborhoodPharmacist”, “FriendlyPharmacist”

“Heisenbergmontana”, “HeisenbergMontna”

“Doritos”, “Doritos_ Doritos ”

“Anonymous-Narcs”, “AnonymousNarcotics”

“swissweed12”, “swissweed”

“DrWhiteTeam”, “DrWhiteInt”

“Kript0x”, “L0rzo”

“hamermike”, “mikehamer”

“books4theunderground”, “books4theug”

“skypeman_ evo”, “skypeman_ accounts”, “skypeman”

“rockyroad1988”, “rockyroad14”, “rocky-shardy”

“dotzvuss”, “dotz.vuss”

“2fargone”, “Tofargone”

“UMIT”, “Lumitrad”

“c63amg”, “c63amgSR”

“Maggots”, “Maggotz”

“coffeeshop”, “UKCoffeeShop”

“LDN-UNDERGROUND”, “LONDON-UNDERGROUND”

“goingpOZtal”, “GoingPostalGroup”

“bonnie-clyde”, “BonnieNClyde”

“TrapHouseBM”, “TrapHouse”

“QueenGreen”, “GreenQueens”

“Zeus”, “Zues”

“R.I.P.”, “RastainPeace”

“swazidoctor888”, “swazibudbud888”

“Alfa& Omega”, “AlfaOmega”

“JipTraVoltaHT”, “Jip-TraVolta”

“MysticRideThroughGalaxy”, “RideTroughGalaxy”

“DMTlovestore”, “DMT-lovestore2”

“Optimum_ Cannabis”, “Optimum_ Cannabinoid”

“CaliBud”, “CB2013”

“chomper”, “chomperlegit”

“AnnKatarinRosenblad”, “AnnKRosenblad”

“DrogMann”, “Mann-Drog”

“MadeinGermany88”, “MadeinGermany”

“WeedGirls”, “WeedGirlz”

“Warlord5000”, “Warlord3000”

“lafl0che”, “lafloche”

“ChemBrothersAU”, “chemicalbrothers”

“Postman-Pot”, “MrPostmanPot”

“indiabenzos_ ib”, “indiabenzos”

“mrblacklabel”, “blacklabel”

“Sunshining”, “sunshines”

“Original_ smileawhile”, “smileawhile”

“Passport_ Connection”, “passconnectio”

“ukmedsnew”, “ukmeds”

“Quixote”, “Quixotic”

“RecConscious”, “ReconnoiterConscious”

“SC_ Connect”, “socalconnect”

“zetaze”, “ZetaOC”

“COCAINECOWBOY-SR”, “Cocaine_ Cowboy”

“Weedsmoker”, “weedsmoker88”

“ntts”, “NOWTHATSTHESTUFF”

“PillsThrillsChills”, “ChemsPillsThrills”

“SUPERMARIO-OFFICIAL”, “SUPERMARIO-EVO”

“SeattlesBestWeed”, “seattlesbestcannabis”

“eJuiceKing”, “eCigJuiceKing”

“chronicbuds”, “chronicbudz”

“GVtobaccoAL”, “gvforsaleal”

Appendix C

List of Subcategories

- **Psychedelics** (47) 4-aco-dmt, lsd, lsb, tma-family, lsa, shrooms, 5-meo-dmt, 4-ho-dipt, other pyschedelics, 4-ho, lsz, lb, amt, bufotenin, samples, 3c-p, dmt-family, mpt, 5-meo-amt, andere, dopr, fluff, dpt, aet, lucy, truffles other, 2c, proscaline, 4-ho-met, lysergamides, doses, acid, mipt-family, nmbome, nbmd, tryptamine, methallylescaline, 4-propo, 3c, psychedelics, 2c-family, 5-meo-family, met, ibogain, others, 2c-p, 4-aco-det, 2c family, 2c-b, 2c-c, 4-aco, 2c-d, 2c-e, amt-family, 2c-i, 4-aco-met, 3c-family, 4-meo-family, 5-meo-mipt, al-lad, mushrooms, spores, microdots, det-family, tabs, psychedelic, mimosa hostilis, cultures, schrooms, enthogens, dipt, 4-ho-family, mescaline, allylescaline, mushrooms - grow, mescaline, det, escaline, dmt, doc, 4-aco-family, 5-meo-dalt, 5-meo-dipt, 4-aco 5-meo, dipt-family, psilocybin mushrooms, tab, 4-aco family, trip, nb, 5-ho-family, salvia, other, truffles, blotter, hot-7, allylescaline, 2c-t-7, dalt, ayahuasca, 4-ho-mipt, 2c-t-2, 5-meo, nbome, mipt, pipt, doi, don, liquid, dom, dob, nboh, dox, muscimol, dot, 2c-i-nbome, 25i-nbome, cimbi-5, candy
- **Prescription** (47) anti-depressant, general, antidepressants, sexual enhancement, analgesics, sleep aids, hydro, other, samples, general health, prescription, prescription, relaxants, SSRIs, MAOIs
- **Opioids** (44) china white, extreme pain meds, tramadol, methylmorphine, oxycodone, loperamide, meptazinol, pentazocine, dezocine, black tar heroin, hydromophone, fentanyl, sufentanil, samples, tilidine, stamp bag, morphin, nicomorphine, diphenoxylate, express, butorphanol, levacetylmethadol, levomethorphan, nalmefene, fentanyl/other, pepap, morphone, sildenafil citrate, dextro propoxyphene,

hydro/oxycodone, oripavine, hydro/ oxymorphone, kratom, oxycodon, pills, ohmefentanyl, dextromoramide, oxygesic, nalbuphine, dihydrocodeinone, ah-7921, remifentanil, substitutes, heroin, others, prodinenisentil, prescription, levorphanol, a-methylfentanyl, oxycontin, phenazocine, heroin #3, etorphine, allylprodine, naloxone, desmethylprodine, instant release, stamp, naltrexone, buprenorphone, dihydroetorphine, tapentadol, oxymorphone, dimorphone, buprenorphine, alfentanil, opium, oxynorm, opioids, morphium, other opioids, meperidine, heroin #4, ketobemidone, coedine, hydrocodone, bezitramide, hydromorphone, analgesics, panda, dipipanone, other, suboxone, paramorphine, lefetamine, piritramide, hydrocone, carfentanyl, dihydrocodeine, heroine, morphine, codeine, methadone, meperidine, Roxicodone, Vicodin, pethidine, norco, buprenorphine, Demerol, fentanyl

- **Steroids** (36) antagonists, agonists, fluoxymesterone, drostanolone, aromatase inhibitors, human grow hormones, winstrol, other, hcg, metabolism, 25ng, human growth hormones, methandrostenolone, clenbuterol, stanozolol, anavar, mesterelone, anabolic steroids, other steroids, oxandrolone, injectable, oral, testosterone, 20mg
- **Ecstasy** (46) pressed pills, ecstasy/mdma, methylone, mda, 5-it, ecstasy, ap-family, tan, pills, 4-mec/other, md-family, mandy, methylone, other, xtc, butylone, m1, white, 4-mec, pill, brown, capsules, samples, mpa, molly, ethylone, pentedrone, pentylone, tan mdma, mda/mdxx, others, other ecstasy, sass, bk-mdma, mdma, midai, crystal, mdai, methlone, ecstasy pills
- **Cannabis** (55) pre-rolled, trim, aaa, outdoor, oil polen, indica, cuttings, skywalker, satvia, indoor, marijuana, wax-pollen, seed, medical marijuana, other cannabis, weed), edibles, synthetic, kush, co2, cbd, purple, shake/trim, hybrid, oil and pollen, cannabis, other, samples, hashish, pot, wax, synthetics, oil, hash, oils, shatter, vegan, weed, shake/trim/kief, concentrate, others, clones, topicals, seeds/ synthetics, flowers, grow, bud, edible, syntetics, oil and polen, pre-rolled/trims, extracts, oil pollen, true landrace, steroids, og, pre-rolled joint, seeds, concentrates, shatter, high-grade, resin, carramello, pressed
- **Benzos** (32) flutoprazepam , dmcm , chlordiazepoxide , clotiazepam , flunitrazepam , triazolam , premazepam , eszopiclone , halazepam , medazepam , nitrazepam, clorazepate , b1, 7.5mg, cloxazolam , delorazepam , etizolam, triangles, bars, diazepam, nordazepam , lorazepam , usa, ethyl loflazepate , blotter, bretazenil , alprazolam, clobazam , other, tetrazepam , klonopin, zolpidem , lormetazepam , temazepam , midazolam , xanax, clonazepam, lorazepam, oxazepam , ativan, powder, pinazepam , mix, cinolazepam , zopiclone, ketazolam , brotizolam , pyrazolam , valium, estazolam , flumazenil , loprazolam , flurazepam , bromazepam , quazepam , thienodiazepine, benzos, 20mg, nimetazepam , prazepam , phenazepam , zaleplon , 2mg

- **Dissociatives** (39) poppers, salvinorin a, 8a-pdhq, salvinorin b, dioscorea, phencyclamine, enadoline, other dissociatives, metaphit, 4-meo-pcp, pills, k, ketamine, u-50488, nefa, dextrorphan, methoxetamine, pcp, spiradoline, dizocilpine, remacemide, rolicyclidine, phencyclidine, other, jenkem, nitrous oxide, hz-2, gacyclidine, 3-meo-pcp, methoxyketamine, ghb, dexoxadrol, tiletamine, mxe, n-ethylnorketamine, 2-mdp, tifluado, dextromethorphan, dxm, diethyl ether, scopolamine, esketamine, tenocyclidine, eticyclidine, gbl
- **Stimulants** (50) cathinones, converta, "fas", moppp, ethylphenidate, mephedrone, 5-apb,6-apb,a-pvp-cocaine, d2pm, 5-apb, fishscale, meth, 6-apb, speed, cathinone, cocaine, ritalin, 3,4dmmc, fmas, eightball, crystal meth, caffeine, benzedrine, other, samples, mdpv, amphetamines, vyvanse, khat, coca leaves, methamphetamine, mdppp, ghb, dextroamphetamine, cocoine, coke, 4-emc, pentedrone, dexedrine, ephedrine, "fmcs", amphetamine, a-pvp, prescription, cocaine meth, crack, dimethocaine, 2-dpmp, meow meow, kokcain, adderall, speed paste, others, other stimulants
- **Other Drugs** (39) weight loss, pharmacy, precursors, barbitures, research chemicals, paraphernalia, substances, rcs, sample, alcohol related products, samples, cultures, inhalants, nootropics, barbituates, misc, barbiturates, entheogens, wholesale, phenethylamine, honey oil, intoxicants, organic, tobacco, supplements, pressed, whole sale, antidotes

Appendix D

List of Common Words

swazi, alice, cooki, evendor, weedd, castle, custom, termin, ishop, worldwide, middle, viking, apples, ernotratedyet, limit, andro, smack, smoker, anders, anman, pshop, egood, estor, ranger, lling, deman, paris, phone, cherry, uncle, porter, felix, thepo, grams, stoner, allen, reeze, ersto, llion, onald, ndrugs, fruit, rasta, walter, times, econnect, esand, chroni, france, edeal, contr, armer, chard, goods, jimmy, ssion, blaze, ington, oming, highquality, solutions, bestb, edeli, thegreen, stoned, theke, opiate, terda, thegreat, corne, color, ration, eworld, guides, kauppa, blackha, digitalp, aptain, amine, hustle, lking, medicin, special, roman, spharma, ocket, cious, krypt, anonym, electr, multi, econn, christ, strai, rdrugs, depot, bestd, ershop, freed, dutchp, steve, deale, sunny, sterm, speedy, tting, budma, dutchm, horse, gandalf, mster, yourm, johnson, tender, onymo, lower, stero, ookie, weedy, blackl, apothe, ofweed, speci, realp, dyman, andr, ydrug, label, concentrate, itali, humboldt, thepr, emerald, owers, cannabisc, herba, richard, nurse, arter, erati, offic, thebig, sting, ucker, iking, erthe, peman, ground, ctive, silen, rando, baron, lights, xpert, parad, dicat, redbull, mafia, thego, rance, ycrew, estin, cross, edrug, secrets, smooth, sugar, nextday, calic, george, sells, ideal, orthe, sauce, blues, devil, aussies, jesus, shoppe, bobby, rgreen, siness, antas, american, walker, buyer, group, nline, ehouse, double, nigga, original, delivery, sfinest, psycho, theb, mountain, deliver, baker, tweed, supplies, check, eagle, weedg, stand, silky, andym, besto, anotratedyet, npharma, etrad, pirat, thern, official, freak, jones, charles, glass, concentr, friendly, beard, bigbo, girls, disco, flying, andma, dolla, hards, weede, grower, carte, nshop, franc, dutchquality, demon, california, ydrugs, darkn, tothe, dnotratedyet, travel, lotus, ehead, cream, company, norway, theblack, ipper, sking, three, thefa, ealth, discount, stick, etime, smile, perso, maker, terdam, fraud, ocaine, kweed, monster, cheese, wonder, highs, supplier, trust, thebest, tamin, ebear, theca, turtle, santa, products, daddy, crazy, everything, wholesale, pablo, minal, arder, blackb, topsh, ution, hello, greenl, brain, ereal, rabbi, crypto, sleep, kshop, amazon, blanc, emical, solut, rwhite,

darkne, ports, ittle, prod, cardi, cocai, sterd, yking, steal, euros, narco, legal, america, lover, theon, genera, emoney, anger, grand, ester, westcoast, anony, theba, lland, earth, ncent, adder, sweed, mario, centra, anonymous, inter, germany, darks, rshop, brother, lemon, sters, sensi, perfect, onkey, liber, thedark, ebooks, astro, herbs, snake, study, delic, chemicals, tnotratedyet, general, chron, thedo, shine, dshop, lands, train, paper, books, steroid, estic, buddy, thegr, themo, rypto, family, marijuana, thereal, morph, therea, china, theco, cation, andal, prime, epharma, honest, carde, dking, senberg, chill, crack, carder, cloud, planet, organic, silver, panda, cannabi, estco, killer, thebl, organi, warrior, frost, unter, billy, roger, sport, stcoast, kitty, underground, ygirl, carding, johnny, nnotratedyet, ntain, project, thebe, theda, storm, olution, eshop, thebo, brown, feelgood, isell, ather, ebook, nster, sdirect, ackha, expres, acker, mushroom, royal, nking, secret, south, prince, drugstore, ymous, crystal, herbal, purple, account, finest, ander, order, coast, trader, source, right, erson, liver, thegre, hquality, phoenix, darknet, ninja, shrooms, rnotratedyet, adderall, enotratedyet, agora, ofthe, friend, heaven, organ, class, guide, juice, deral, caine, tripp, supplie, brothers, gener, sunshine, shard, medica, peter, edrugs, clean, french, masters, chems, peace, johns, chris, dragon, frank, eweed, paypal, onion, online, business, medical, agent, street, press, machine, grade, every, thech, conne, dreams, wizard, powder, australia, pirate, shaman, icals, dutchd, atter, armacy, cartel, hydro, hacker, seeds, trans, cking, monkey, cards, ights, premium, thema, range, ender, thers, amster, nation, chemist, thing, alien, angel, import, annabis, rious, sshop, deals, lucky, product, nabis, elite, cocain, charlie, ronic, ealer, great, canadian, onest, iller, north, round, other, apple, quick, shadow, garden, psych, amsterdam, rabbit, digital, greens, sweet, queen, ganja, trated, notra, eking, ghost, little, benzo, supers, smith, golden, aking, power, stealth, sales, enberg, marke, james, charl, mister, weeds, party, there, candy, service, diamond, holland, stuff, chronic, dealer, harma, canada, chemical, armac, swiss, oking, domestic, cheap, captain, snotratedyet, inger, ection, ister, tions, stone, rated, fresh, alpha, space, count, under, ality, cannab, chemi, xanax, canad, pharmacy, ction, night, medic, magic, world, trade, aster, light, pharmac, cocaine, house, suppl, molly, kings, smoke, eller, shroom, dream, supply, connection, german, market, direct, happy, seller, aussie, pills, money, doctor, xpress, ation, speed, trate, vendor, master, cannabis, erman, express, super, quality, canna, connect, store, white, pharm, black, pharma, dutch, green, notratedyet, drugs

Appendix E

Reputation

1776

The vendor reputation value provided by 1776 was used to rank vendors.

Abraxas

The feedback rating provided by Abraxas was used to rank vendors.

Agora

The average rating across each vendors reviews was used to rank the vendors.

Alpaca

The rating was taken to be the vendor rating given.

Alphabay

The rating was taken to be the vendor level given with a higher number assumed to be better.

Amazondark

No feedback data was available for vendors or product listings. Instead, vendors were given a rating of 0 if they did not pay for their vendor account and 1 otherwise.

Anarchia

The rating was taken to be the vendor rating value given.

Andromeda

The rating was calculated as the proportion of positive reviews out of the total number of reviews.

Area 51

The rating was taken to be the average rating provided.

Black Services Market

The rating was taken to be the vendor rating value given.

Bloomsfield

The rating was taken to be the vendor rating value given.

Blue Sky

The rating was taken to be the vendor rating value given.

Breaking Bad

The rating was taken to be the vendor rating value given measured as a percentage.

Cannabis Road 2

The rating was taken to be the vendor rating value given.

Cannabis Road 3

The vendor rating (the number of points they have) was used to rank the vendors.

Cantina

Vendor profiles did not contain feedback or vendor rating information and so vendors were ranked using the average rating giving to their products.

Cloudnine

The rating was taken to be the positive rating divided by the total rating.

Cryptomarket

The rating is the number of positive reviews divided by the total number of reviews.

Darkbay

Darkbay provided multiple types of reputation information the number of reviews, the positive, neutral, and negative ratings, feedback, and the seller level. Because the number of positive, negative and neutral ratings was more granular than the seller level and better distinguished between vendors, these values were used for the vendor ranking. They were combined by dividing the number of postive ratings by the total number of ratings.

Darklist

For Darklist, only feedback comments were available. The classifier was used to calculate the reputation score.

Dark Net Heroes

Each vendor was assigned a member level, this was used to rank vendors.

Deepzon

The vendor rating was used to rank vendors.

Diabolus

The proportion of positive ratings out of the total number of positive and negative ratings was used to rank vendors.

Dogeroad

The vendor pages contained no review or rating data so the vendor rating value on each listing page was used rank vendors instead.

Dreammarket

The vendor rating was used to rank vendors.

Drugslist

The average of the product quality, delivery speed, and communication ratings was used to rank vendors.

Eastindia Company

The vendor level, as opposed to rating, was used for ranking as this was of a greater level of granularity.

Evolution

The Evolution site provided several types of vendor information, including review data, a vendor rating, the percentage of positive feedback and a user ranking. Not all of the users had feedback and the vendor rating and percentage of positive feedback did not offer much distinction between vendors. The user ranking system provided a useful metric and was used to produce the vendor rankings. Two different systems were used. To begin with, vendors were ranking on a scale of Freshman, Sophomore, Junior, Senior, Premium, Advanced, Expert, Master, Grandmaster, Godlike. The system then switched to levels 1,2,3,4,5. There is a clear distinction in the dates when each system was used.

Freebay

The vendors were assigned a Seller Level, which was used to rank the vendors.

Freemarket

No feedback information was provided. **Greyroad**

The vendor ratings were used to rank vendors.

Haven

The vendor rating was presented as the percentage of feedback which was positive, this was used to calculated the rankings.

Horizon

No information was available.

Hydra

The vendor rating was used to rank vendors.

Ironclad

The vendor rating was presented as the percentage of feedback which was positive, this was used to calculated the rankings.

Kiss

The vendor rating was used to rank vendors.

Middleearth

The vendors were ranked using the number of positive ratings divided by the total number of ratings because every vendor had this information recorded for their account, this was not the case for other types of vendor information, for example the vendor level.

Nucleus

Two measures were available the vendor level and the number of reviews given 1 to 5 stars. Every vendor had the latter recorded but not every vendor had a user a level and so a rating was constructed from the average number of stars (total number of stars received divided by number of reviews) and used to rank vendors.

Onionshop

The vendor rating was used to rank vendors.

Outlaw Market

The average star rating of the vendors comments was used to rank vendors.

Oxygen

Because none of the vendors had reviews on the Oxygen marketplace, the vendor statistics were used to rank vendors. These statistics contained information on the number of sales, as well as the number and proportion of successful sales, disputed sales, and sales in progress. The proportion of successful sales was used.

Panacea

Where available, the vendor rating was used to rank vendors. For the few vendors who did not have a vendor rating, the number of positive ratings minus the number of negative ratings was used as this produced the same value for all vendors for whom both types of information was available.

Pandora

The vendor rating (a score out of 5) was used to rank the vendors.

Pigeon Market

The vendors were ranked using the proportion of positive ratings.

Pirate Market

The vendor rating was used to rank vendors. Where it was not available (due to a collection error that occurred on two occasions) the previous months rating was used.

Poseidon

The vendor rating (measured out of five) was used to rank vendors.

Sheep

The rating was taken to be the overall vendor reputation measured as a percentage.

Silk Road 2

The average review score (out of 5) was used to rank the vendors for the scrapes where it was available, where not the rating was computed using the review classifier.

Silk Road Reloaded

The average value across each of the vendor ratings was used to rank the vendors.

Silk Street

Whilst vendor profiles did store review data on SilkStreet, no vendor reviews were made and so this information could not be collected. Instead, the vendor rating stored on listing pages was used to rank vendors.

the Majestic Garden

There is no review data on vendor or listing pages.

the Marketplace

The average ranking (valued as a percentage) was used to compute the vendor rankings.

the Real Deal

The vendor rating was used to compute the vendor rankings.

Tochka

The vendor rating was used to compute the vendor rankings. However, only a very small proportion of vendors had this information available.

Tom

Review data was available in the form of comments and ratings, the average rating across all reviews for each vendor was taken and used the rank the vendors.

Topix 2

The average vendor rating from the review data on the vendor profile was used to rank vendors.

Tor Escrow

The user rating was used to rank vendors. Some users did not have a percentage rating, instead they were labelled as new, initiate, credible, and established vendors. New and initiate vendors were given rating values of 0, credible vendors 50 and established vendors 100.

Tor Market

The user rating was used to rank vendors.

Torbay

The vendor rating was used to rank vendors.

Torbazaar

The average rating was used to rank the vendors. **Tortuga2**

The user rating was used to rank vendors. Some of the data was collected incorrectly, e.g. had location data instead of vendor rating, these vendors were given the rating from the previous month.

Underground Market

The vendor rating was used to rank the vendors.

Utopia

The vendor rating was used to rank the vendors.

White Rabbit

The vendor rating was used to rank the vendors, where this was not available, the average rating on the vendors listing was used instead.

Appendix F

Timeline

April 2010 Mephedrone classified as class B (DrugWise (2016))

27 January 2011 advert/mention in `shroomery.org` (probably by RU) (Bearman and Hanuka (2015a))

March 2011 Bad Apple attack on Tor shown to be successful, attack identified users of BitTorrent but could also be applied to other vulnerable sites (Blond et al. (2011))

1 June 2011 Gawker ran a story on *Silk Road* leading to a surge in popularity of the site (Chen (2011))

June 2011 two U.S. senators publish a letter calling for the banning of bitcoins which references *Silk Road* (Manchin (2011))

February 2012 announcement of name Dread Pirate Roberts (Bearman and Hanuka (2015b))

1 September, 2011 Stacy Litz, a buyer (of MDMA and LSD) on *Silk Road*, arrested in Philadelphia (Branwen (2017))

16 April 2012, 8 people arrested and DNM *The Farmers Market* (formally Adamflowers) is shut down in operation AdamBomb, a 2 year multinational operation (DEA (2012))

April 2012 Methoxetamine (like Ketamine) is given a TCDO banning sale (but not possession) (DrugWise (2016))

September 2012 *RAMP* launched

November 2012 Methoxetamine and some new synthetic cannabinoids (incl. "Black Mamba") are classified as class B drugs (DrugWise (2016))

17 January 2013 Curtis Green is arrested in an FBI sting (Mac (2013b))

3 March 2013 discussion on Reddit about a worrying security warning on *Silk Road* (Reddit User (2013a))

26 March 2013 discussion on Reddit about whether or not *Silk Road* revealed its IP address (Reddit User (2013b))

3 April 2013 Ross Ulbricht states in his logs that scams were becoming a problem (Bearman and Hanuka (2015b))

2 May 2013 Ross Ulbricht states in his logs that attacks (which he has to pay protection money for) are causing Tor to crash. Around this time, the site was shut down for a week by attackers (Bearman and Hanuka (2015b))

26 May 2013 RU states in his logs that he accidentally leaked the IP twice whilst trying to move the forum to multi.onion config (Bearman and Hanuka (2015b))

June 2013 NBome (related to 2Cl) and Benzo Fury (related to ecstasy) given temporary classification drug orders and khat is classified as a class C (DrugWise (2016))

4 August 2013 founder of Freedom Hosting arrested Buterin (2013)

14 August 2013 Forbes interview with Dread Pirate Roberts which describes some security failures but also proselytises some of the *Silk Road* philosophy. The interview links to `silkroadlink.com` which gave instructions and advice on how to access *Silk Road* Greenberg (2013b)

August 2013 Tormail goes offline after an FBI raid on FreedomHosting (Wikipedia (2017))

5 September 2013 an article about tracing BTCs to *Silk Road* published in Forbes (Greenberg (2013a))

18 September 2013 an article about the arrest of Adam Bunger, accused of selling guns on *Black Market Reloaded* published in the Daily Dot (O'Neill (2013)). This was followed by quite a lot of discussion about his arrest and the deletion of some comments that risked providing more information for a police case against him.

20 September 2013 *Atlantis* closes due to "security reasons" (Market (2013))

27 September 2013 criminal complaint signed on Ross Ulbricht (U.S. District Court for the Southern District of New York (2013))

3 October 2013 FBI seizure and closure of *Silk Road* (Van Buskirk et al. (2014))

October 2013 FBI arrests Ross Ulbricht and obtains access to *Silk Road* posting a notification of seizure onto the *Silk Road* site (Jeffries (2013); Bearman and Hanuka (2015b)). `freeross.org` website set up to free Ross Ulbricht

17 October 2013 *Black Market Reloaded* announces that it is shutting because of a security vulnerability and then doesn't, claiming the vulnerability was not as bad as first thought (Mac (2013a))

21 October 2013 *Pandora* established (Branwen (2013))

28 October 2013 Curtis Green criminal complaint signed (Mac (2013b))

5 November 2013 *Dream Market* established (Branwen (2013))

6 November 2013 *Silk Road 2.0* launched (Greenberg (2013c))

7 November 2013 Curtis Green criminal complaint unsealed (Mac (2013b))

15 November, 2013 *Dreammarket* established (Branwen (2013))

19 November 2013 *Piratemarket* established (Branwen (2013))

25 November 2013 due to new or more rigorous screening methods employed by Australian customs, an unusually large number of packages to Australia are not being received. This led to lots of discussion from unhappy Australian customers on forums and negative reviews which prompted vendors to stop shipping to Australia or to change their terms and conditions, refund policy, or expected shipping time when shipping to Australia. (Date attributed to one such vendor post specifically about their new Australian refund policy.) (Martin (2014))

28 November 2013 *the Marketplace* established (Branwen (2013))

1 December 2013 the marketplace *Sheep* closed, reportedly in response to the exploitation of security vulnerability that allowed a seller on the site to steal $6 million in Bitcoin. As a result of the closure, the admin has gained access to $44million in bitcoin that was tied to the site. In response, *Black Market Reloaded* claim that they cannot handle the influx of people that this would lead to and so blocked new applications, suspended inactive accounts and closed the site for a few days (Greenberg (2013d))

2 December 2013 *Black Services Market* and *Freebay Market* 2013 established (Branwen (2013))

3 December 2013 *Agora* marketplace established DarkNetMarkets.org (2013) *Blue Sky* established (Branwen (2013))

9 December 2013 *Silk Road 2.0*, *Pandora Openmarket*, and *Tormarket* suffer DDoS attacks. *Silk Road 2.0* claims the attack was launched by *Tormarket* and so the DDoS of *Tormarket* is maybe a retaliation (Digital Citizens Alliance (2013)) *Silk Road 2.0* was infiltrated by an FBI agent who ended up on the payroll and was eventually able to access the server address. The site was linked to 26 year old software developer Blake Benthall because the server was controlled by someone using the email address blake@benthall.co.uk, the same email address used by Blake Benthall on his public sites (Mullin (2014))

14 December 2013 Dread Pirate Roberts claims to have stolen the entire *Tormarket* database (private messages, addresses, purchases, statistics) and also claims this was to test *Tormarket's* security. Within a fortnight, *Tormarket* closes taking all user bitcoins with them (Digital Citizens Alliance (2014))

16 December 2013 *Tortuga2* market established (Branwen (2013))

18 December 2013 *Torbay Market* established (Branwen (2013))

21 December 2013 3 site administrators' arrests confirmed (DeepDotWeb (2013)). Not long after this Dread Pirate Roberts disappears and "Defcon" takes over as temporary admin (Digital Citizens Alliance (2013))

23 December 2013 *White Rabbit Market* established (Branwen (2013))

29 December 2013 *Outlaw Market* established (DarkNetMarkets.org (2013))

4 January 2014 *Grey Road* established (Branwen (2013))

5 January 2014 *Tortuga2* closed down (Branwen (2013))

14 January 2014 *Evolution Market* established (Branwen (2013))

18 January 2014 *Dogeroad*, *Drugslist Market* established (Branwen (2013))

20 January 2014 *Cantina Market* established (Branwen (2013))

26 January 2014 *Torbazaar* established (Branwen (2013))

30 January 2014 *Tochka Free Market* established (DarkNetMarkets.org (2013))

1 February 2014 *Black Services Market*, *White Rabbit*, and *Grey Road* closed down (Branwen (2013))

2 February 2014 *Tor Escrow* established (Branwen (2013))

3 February 2014 Schumer writes another open letter about online market places and the fact that they still exist despite the closure of *Silk Road* (Greenberg (2014d)). This letter was discussed on `silkroad.org` (Silk Road Drugs (2014)). *Utopia* marketplace launches (Greenberg (2014a))

7 February 2014 *Cantina* closed down (Branwen (2013))

9 February 2014 *Black Goblin Market* deanonymised by Gwern (Reddit User (2014b))

11 February 2014 *Cloud 9* marketplace established, *Utopia* closed down (Branwen (2013))

12 February 2014 *Utopia* closed down by Dutch authorities (Greenberg (2014a)) 5 arrested (BBC (2014b))

13 February 2014 *Silk Road 2.0* announces that it has been hacked, with an estimated loss of around $27million, citing "transaction malleability" as the cause. However, many *Silk Road 2.0* users reject the explanation and blame either administration incompetence or a deliberate scam on behalf of the *Silk Road 2.0* administration team (Greenberg (2014e); DeepDotWeb (2014e))

15 February 2014 *Silk Road 2.0*, *Budster* and *Agora* go down within 36 hours *Silk Road 2.0* because of a hack, *Budster* in an exit scam, and *Agora* because of an overload from *Silk Road 2.0* users (Branwen (2013))

19 February 2014 *Dark Net Nation* established (Branwen (2013))

24 February 2014 *Nucleus* established (Branwen (2013))

28 February 2014 *Freebay* and *Drugslist* closed down (Branwen (2013))

12 March 2014 *Hansa* IP address leaked (DeepDotWeb (2014a))

13 March 2014 *Dogeroad* established (Branwen (2013))

19 March 2014 *Pandora* also announces it has been hacked with half its Bitcoin ($250,000) stolen (DeepDotWeb (2014c))

20 May 2014 *Pandora* hacked losing half of escrow ($250,000) (DeepDotWeb (2014c))

23 March 2014 *Red Sun* marketplace hacked by the Avid who advertised this on Reddit (Reddit User (2014e))

24 March 2014 *EXXTACY* opens and closes in a day (Reddit User (2014f))

25 March 2014 *Topix 2* established (Branwen (2013))

27 March 2014 *Hydra* established (Branwen (2013))

28 March 2014 *Cannabis Road 2* established (Branwen (2013))

5 April 2014 *Andromeda* established (Branwen (2013))

14 April 2014 *Pigeon Market* established (Branwen (2013))

19 April 2014 *1776* established, *Tor Escrow* closes down (Branwen (2013))

20 April 2014 *Alpaca* established, *Torbay* closes down (Branwen (2013))

23 April 2014 *Underground Market* established (Branwen (2013))

April 2014 Heartbleed OpenSSL bug takes Tor offline for several days whilst admin remove vulnerability (Synopsis Inc (2014)). Grams search engine launched allowing users of DNMs to search for specific products across multiple markets (Reddit User (2014c))

1 May 2014 *Darkbay* closes down (Branwen (2013))

7 May 2014 *Pigeon Marketplace* closes down (Branwen (2013))

May 2014 *the Majestic Garden* marketplace is established (Tarquin (2014)) with some teething problems (DeepDotWeb (2014b)). *Alpaca* hacked and information "deleted" (Reddit User (2014a)).

9 May 2014 *Armory* accepts Darkcoin (Reddit User (2014h))

10 May 2014 *Tom* established (Branwen (2013))

14 May 2014 *Deepzon* established (Branwen (2013))

18 May 2014 *Onionshop* established (Branwen (2013))

23 May 2014 *Horizon* established (Branwen (2013))

20 May 2014 *Area 51* established (Branwen (2013))

17 June 2014 *Underground Market* closes down (Branwen (2013))

22 June 2014 *Middle Earth Marketplace* established (DarkNetMarkets.org (2013))

June 2014 "Ketamine reclassified from Class C to Class B in response to concerns about damage to the bladder from long term use. A number of substances are classified, including NBOMe and related compounds which are now Class A, and 'Benzo Fury' and related Benzofuran compounds which are Class B. Lisdexamphetamine, a medicine which converts into amphetamine in the body, is classified as Class B. Tramadol, an opioid painkiller, is classified as Class C, as are Zaleplon and Zopiclone, which are sedatives similar to the already-classified Zolpidem." (DrugWise (2016))

8 July 2014 *Horizon* closes down (Branwen (2013))

13 July 2014 *Deepzon* closes down (Branwen (2013))

4 August 2014 *Silk Street* closes down (Branwen (2013))

15 August, 2014 *Piratemarket* closes down (Branwen (2013))

19 August – 1 September, 2014 *Pandora* administrators locks vendor accounts and stops withdrawals (Branwen (2013))

25 August 2014 *Cannabis Road 2* closes down (Branwen (2013))

20 September 2014 *Onion Market* hacked (Reddit User (2014g))

27 September 2014 *Onionshop* closes down (Branwen (2013))

September 2014 *Silk Road* hacked (Branwen (2013))

2 October 2014 *1776* closes down (Branwen (2013))

6 October 2014 *Cannabis Road 3* established (Branwen (2013))

13 October 2014 *Diabolus Market* established (Branwen (2013))

27 October 2014 *Panacea* established (Branwen (2013)). administrators of DNMs doxxed (Reddit User (2014d))

29 October 2014 an article in coindesk discusses research from the University of Luxembourg that shows using Tor and Bitcoin simultaneously can reduce anonymity (Biryukov and Pustogarov (2015); Coindesk (2014)).

5 November 2014 Blake Benthall, the suspected administrator of *Silk Road 2.0*, is arrested (Greenberg (2014c); U.S. Attorney's Office (2014c)) Seperately, two men in Dublin are found in possession of drugs and bitcoins and are arrested (BBC (2014a)). Operation Onymous begins leading to the closure of *Silk Road 2.0*, *Bluesky*, *Torbazaar*, *Cloud 9*, *Topix 2*, *Hydra*, *Alpaca*, *Cannabis Road 3*, *Flugsvamp*, *Black Market*, *Pandora* amongst 414 `.onion` sites were found leading to the shutting down of over a dozen online market places plus a number of money laundering sites (Greenberg (2014c)) the operation is announced on 7 November 2014 (U.S. Attorney's Office (2014a)). A spokesperson from Tor said that it was likely that individual sites were targeted, as opposed to there being a vulnerability in Tor itself (Greenberg (2014b))

7 November 2014 *Tor Market* established (Branwen (2013))

9 November 2014 *the Marketplace* closes down (Branwen (2013))

18 November 2014 *Andromeda* closes down (Branwen (2013))

13 December 2014 *Abraxas Marketplace* established (DarkNetMarkets.org (2013))

18 December 2014 *Tom* closes down (Branwen (2013))

22 December 2014 *Cryptomarket* and *Alphabay* established (DarkNetMarkets.org (2013))

25 December 2014 *Diabolus* closes down (Branwen (2013))

December 2014 *Alphabay* introduces digital contracts (Cox (2016b))

13 January 2015 *Silk Road Reloaded* established (Branwen (2013))

14 January 2015 *Freemarket* established (Branwen (2013))

24 January 2015 *Area 51* closes down (Branwen (2013))

30 January 2015 *Tochka* established (Branwen (2013))

13 February 2015 *Panacea* closes down (Branwen (2013))

14 February 2015 *Cryptomarket* established (Branwen (2013))

16 February 2015 *Freemarket* closes down (Branwen (2013))

19 February 2015 *Kiss* established (Branwen (2013))

3 March 2015 *The Real Deal Market* established (DarkNetMarkets.org (2013))

12 March 2015 Shiny Flakes (vendor) bust reported in `deepdotweb.com` (DeepDotWeb (2015e))

14 March 2014 *Evolution* closes down (Branwen (2013))

17 March 2015 *Evolution* closes in exit scam (Krebs (2015)) documented in a reddit thread (Reddit User (2015d)). *Ironclad* established (Branwen (2013))

25 March 2015 *Agora Marketplace* closes to work on patching vulnerabilities to new attacks which are suspected of taking place (Reddit User (2015a)). *Ironclad* closes down (Branwen (2013))

9 April 2015 *the Real Deal* established (Branwen (2013))

16 April 2015 *Oxygen* established (Branwen (2013)) (alternatively established on 26 April (DarkNetMarkets.org (2013)))

24 April 2015 Six overdoses linked to *Silk Road* (BBC (2015))

28 April 2015 *East India Company* online marketplace established (DarkNetMarkets.org (2013))

5 May 2015 *Haven* established (Branwen (2013))

7 May 2015 *Anarchia* established (Branwen (2013))

21 May 2015 *Kiss* admin claims they have been hacked losing 2 BTC (Reddit User (2015c))

26 May 2015 *Kiss* closes down (Branwen (2013))

27 May 2015 *Dark Net Heroes* established (Branwen (2013))

2 June 2015 *Poisedon Marketplace* doxxed on launch day (Branwen (2013))

6 June 2015 *Haven* closes down (Branwen (2013))

8 June 2015 *Amazondark* established (Branwen (2013)) (alternative date 6 August (Reddit User (2015b)))

19 June 2015 *Poseidon* closes down (Branwen (2013))

15 July 2015 *Agora Market* stops selling weapons (DeepDotWeb (2015a))

9 August 2015 *East India Company* hacked, losing BTC30 (DeepDotWeb (2015c))

27 August 2015 *Oxygen* closes down (Branwen (2013))

6 September 2015 *Agora* closes down (Branwen (2013))

25 October 2015 *Amazondark* disappears (Branwen (2013))

21 December 2015 *Torepublic* database leaked (Redakcja (2015))

23 December 2015 Cyruserv went down taking 7 DNMs with it. The admin wrote an explanation of why they were shutting down and revealed a number of security flaws on the site (DeepDotWeb (2015b))

6 June 2016 a story on `deepdotweb.com` published about the arrest of a buyer on *Silk Road* (DeepDotWeb (2016c))

22 – 28 October 2016 Operation Hyperion (Buntinx (2016))

26 November 2016 19 year old Dark Net vendor, Hedon, who traded in pharmaceutical drugs was arrested in Vienna. It is estimated he made over 18,000 sales from mid-April (DeepDotWeb (2016b))

5 December 2016 3 men arrested on suspicion of selling one kilogram of amphetamines from a vendor site on the dark net, the police stated they had been operating since 2015

(DeepDotWeb (2016e))

6 December 2016 10 people have been arrested in Sweden as part of Operation Hyperion (DeepDotWeb (2016a))

7 December 2016 an article about a man arrested under suspicion of importing heroin, methamphetamine and MDMA found not guilty was published on deepdotweb.com. The man successfully claimed that someone else used his computer to access the dark net and order product. Software to access the dark web was found on his computer and the envelopes containing the drugs were sent to his address, they also matched another envelope found in his home, however all of the envelopes were addressed to a different name. The jury found him not guilty and the judge supported the verdict saying it matched the evidence (DeepDotWeb (2016d))

Appendix G

CDF Plots of Advertised and Observed Listing Values

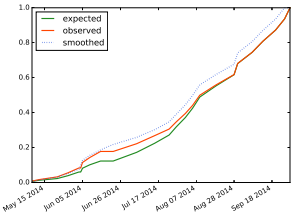

(a) CDF plot of 1776, test statistic 0.120, p-value 0.990

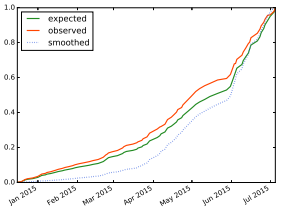

(b) **CDF plot of Abraxas** test statistic 0.065, p-value 0.988

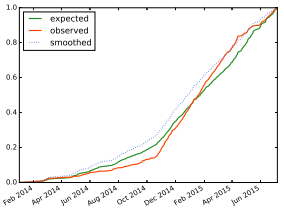

(c) **CDF plot of Agora** test statistic 0.0781, p-value 0.584

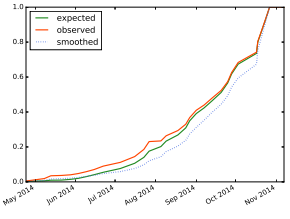

(d) **CDF plot of Alpaca** test statistic 0.161, p-value 0.778

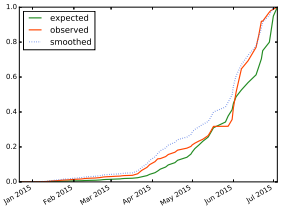

(e) **CDF plot of Alphabay** test statistic 0.167, p-value 0.246

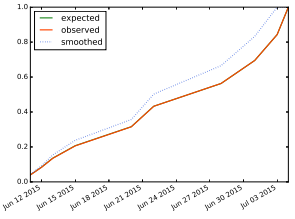

(f) **CDF plot of Amazondark** test statistic 0.100, p-value 1.00

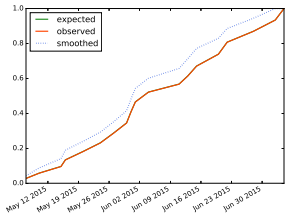

(g) **CDF plot of Anarchia** test statistic 0.0526, p-value 1.00

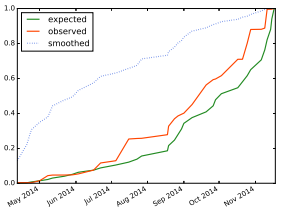

(h) **CDF plot of Andromeda** test statistic 0.162, p-value 0.676

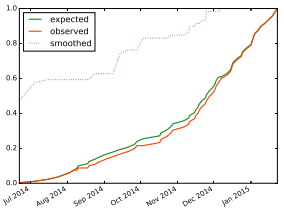

(i) **CDF plot of Area 51** test statistic 0.500, p-value 1.00

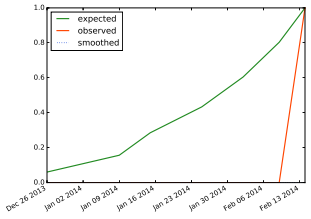

(j) **CDF plot of Black Services Market** test statistic 0.857, p-value 0.00417

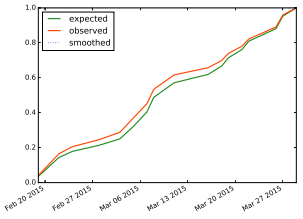

(k) **CDF plot of Bloomsfield** test statistic 0.0526, p-value 1.00

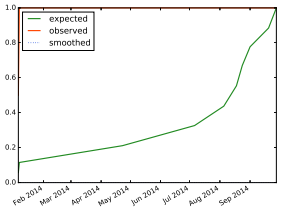

(l) **CDF plot of Blue Sky** test statistic 0.800, p-value 0.00122

FIGURE G.1: **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population.

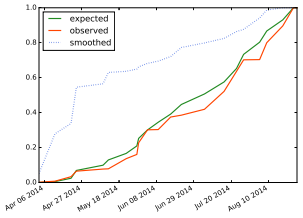

(a) **CDF plot of Cannabis Road 2** test statistic 0.0909, p-value 1.00

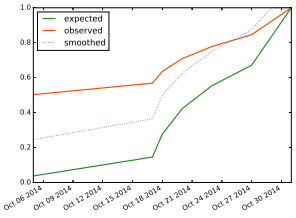

(b) **CDF plot of Cannabis Road 3** test statistic 0.5, p-value 0.188

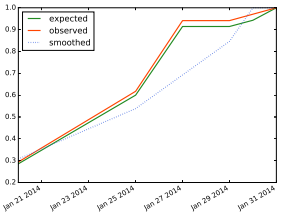

(c) **CDF plot of Cantina** test statistic 0.333, p-value 0.810

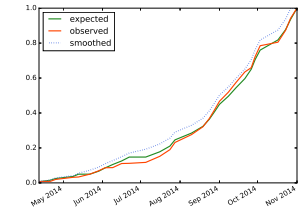

(d) **CDF plot of Cloud 9** test statistic 0.100, p-value 0.997

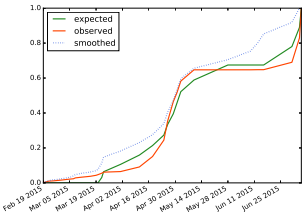

(e) **CDF plot of Cryptomarket** test statistic 0.438, p-value 0.00276

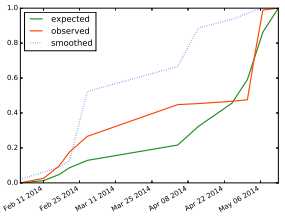

(f) **CDF plot of Darkbay** test statistic 0.167, p-value 0.991

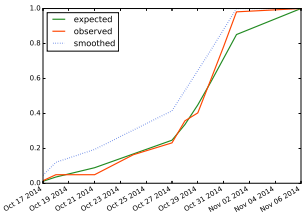

(g) **CDF plot of Diabolus** test statistic 0.0909, p-value 1.00

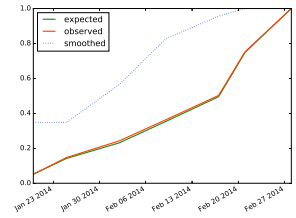

(h) **CDF plot of Dogeroad** test statistic 0.143, p-value 1.00

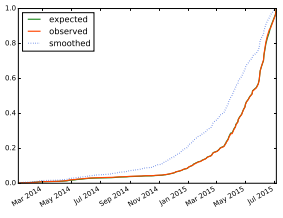

(i) **CDF plot of Dreammarket** test statistic 0.0268, p-value 1.00

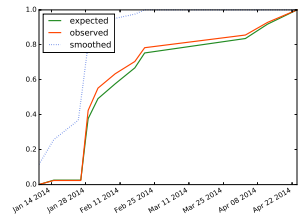

(j) **CDF plot of Drugslist** test statistic 0.214, p-value 0.862

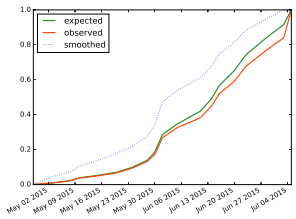

(k) **CDF plot of Eastindia Company** test statistic 0.0455, p-value 1.00

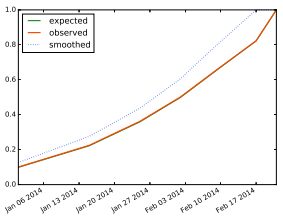

(l) **CDF plot of Freebay** test statistic 0.143, p-value 1.00

FIGURE G.2: **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population.

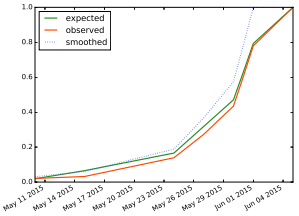

(a) **CDF plot of Haven** test statistic 0.125, p-value 1.00

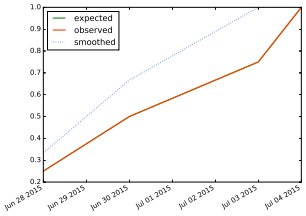

(b) **CDF plot of Horizon** test statistic 0.00, p-value 1.00

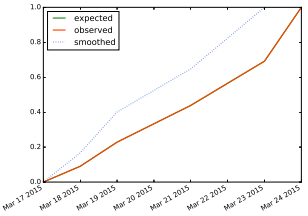

(c) **CDF plot of Ironclad** test statistic 0.167, p-value 1.00

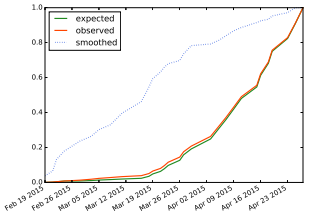

(d) **CDF plot of Kiss** test statistic 1.00, p-value 0.997

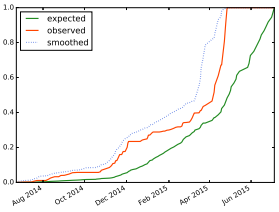

(e) **CDF plot of Middleearth** test statistic 0.280, p-value 0.000142

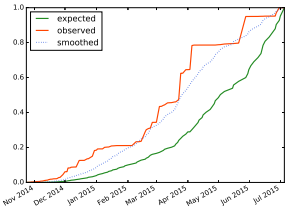

(f) **CDF plot of Nucleus** test statistic 0.261, p-value 0.000603

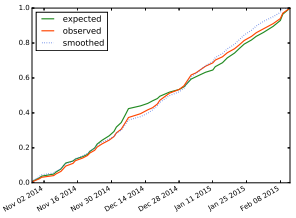

(g) **CDF plot of Panacea** test statistic 0.0556, p-value 1.00

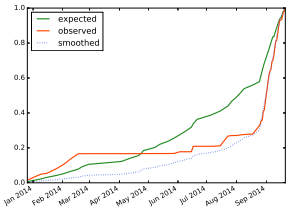

(h) **CDF plot of Pirate Market** test statistic 0.220, p-value 0.0982

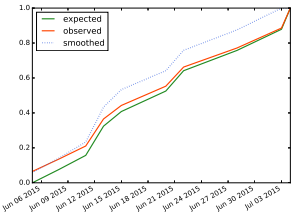

(i) **CDF plot of Poseidon** test statistic 0.0909, p-value 1.00

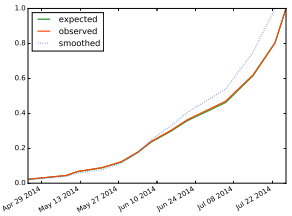

(j) **CDF plot of Silk Road 2** test statistic 0.135, p-value 0.702

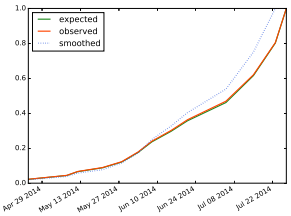

(k) **CDF plot of Silk Street** test statistic 0.0769, p-value 1.00

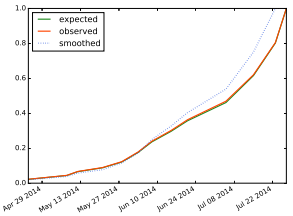

(l) **CDF plot of the Real Deal** test statistic 0.105, p-value 1.00

FIGURE G.3: **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population.

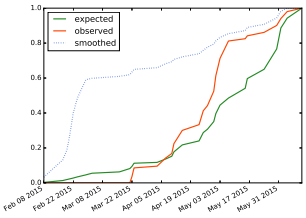

(a) **CDF plot of Tochka** test statistic 0.333, p-value 0.0277

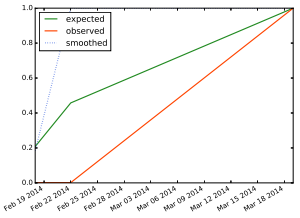

(b) **CDF plot of Tor Escrow** test statistic 0.667, p-value 0.320

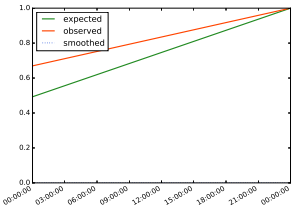

(c) **CDF plot of Tor Market** test statistic 0.5, p-value 0.844

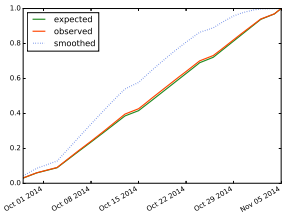

(d) **CDF plot of Topix 2** test statistic 0.0.303, p-value 1.00

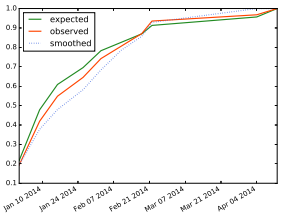

(e) **CDF plot of Torbay** test statistic 0.100, p-value 1.00

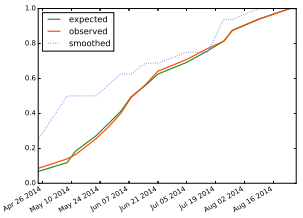

(f) **CDF plot of Underground Market** test statistic 0.0625, p-value 1.00

FIGURE G.4: **CDF Plots and Kolgomorov-Smirnov Statistics** for the Advertised and Observed Number of Listings as well as the CDF Plot of the Smoothed Vendor Population.

Appendix H

ERGM Results

Table H.1: ERGM Results for Network Snapshot on 11/11/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-7.40	$1.61 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	0.000611
Vendor Degree	$4.95 \cdot \exp(-4)$	$1.18 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	0.000611
Trustworthiness	-	-	-	-
Earnings	-	-	-	-
Affordability	-	-	-	-
Diversity	$-2.99 \cdot \exp(-3)$	$3.08 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000609
Age	$1.75 \cdot \exp(-3)$	$3.96 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000612
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-2.80 \cdot \exp(-1)$	1.00	0.780	-
Asia	-	-	-	-
Europe	$-7.61 \cdot \exp(-1)$	$1.09 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	0.000512
Oceania	-	-	-	-
AIC: 9,626, BIC: 9,805				

Table H.2: ERGM Results for Network Snapshot on 18/11/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-7.60	$2.03 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	0.000500
Vendor Degree	$5.50 \cdot \exp(-5)$	$1.42 \cdot \exp(-5)$	0.000101 ***	0.000500
Trustworthiness	$3.37 \cdot \exp(-2)$	$8.14 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000517
Earnings	$-1.19 \cdot \exp(-3)$	$6.72 \cdot \exp(-4)$	0.0755 .	-
Affordability	$3.60 \cdot \exp(-4)$	$8.20 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	0.000500
Diversity	$1.92 \cdot \exp(-3)$	$3.42 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000501
Age	$9.88 \cdot \exp(-4)$	$4.52 \cdot \exp(-4)$	0.0286 *	-
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.07	1.01	0.287	-
Asia	-	-	-	-
Europe	-2.20	$1.40 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$5.54 \cdot \exp(-5)$
Oceania	-	-	-	-
AIC: 9,890, BIC: 10,076				

TABLE H.3: ERGM Results for Network Snapshot on 25/11/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-8.29	$2.12 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	0.000251
Vendor Degree	$-4.49 \cdot \exp(-4)$	$2.56 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	0.000251
Trustworthiness	$1.59 \cdot \exp(-2)$	$3.95 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000255
Earnings	$1.78 \cdot \exp(-3)$	$1.81 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000251
Affordability	$2.53 \cdot \exp(-4)$	$9.64 \cdot \exp(-5)$	0.00885 **	-
Diversity	$2.62 \cdot \exp(-4)$	$2.98 \cdot \exp(-4)$	0.378	-
Age	$-8.62 \cdot \exp(-5)$	$4.35 \cdot \exp(-4)$	0.843	-
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$1.48 \cdot \exp(-1)$	1.01	0.884	-
Asia	-	-	-	-
Europe	-1.49	$1.30 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$5.66 \cdot \exp(-5)$
Oceania	-	-	-	-
AIC: 11,448, BIC: 11,639				

TABLE H.4: ERGM Results for Network Snapshot on 02/12/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-9.21	$9.61 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	0.000100
Vendor Degree	$-2.39 \cdot \exp(-4)$	$2.02 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	0.000100
Trustworthiness	$7.83 \cdot \exp(-3)$	$1.99 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000101
Earnings	$9.17 \cdot \exp(-4)$	$3.72 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	0.000100
Affordability	$3.37 \cdot \exp(-4)$	$4.74 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	0.000100
Diversity	$5.58 \cdot \exp(-5)$	$1.26 \cdot \exp(-4)$	0.658	-
Age	$2.32 \cdot \exp(-3)$	$1.75 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	0.000100
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-9.22	$4.97 \cdot \exp(1)$	0.853	-
Asia	-	-	-	-
Europe	$-4.47 \cdot \exp(-1)$	$5.23 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$6.40 \cdot \exp(-5)$
Oceania	-	-	-	-
AIC: 52,387, BIC: 52,587				

TABLE H.5: ERGM Results for Network Snapshot on 09/12/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-9.75	$1.12 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$5.83 \cdot \exp(-5)$
Vendor Degree	$-2.61 \cdot \exp(-4)$	$1.48 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.83 \cdot \exp(-5)$
Trustworthiness	$3.27 \cdot \exp(-3)$	$5.98 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.85 \cdot \exp(-5)$
Earnings	$2.02 \cdot \exp(-4)$	$1.01 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.83 \cdot \exp(-5)$
Affordability	$7.10 \cdot \exp(-4)$	$4.03 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.83 \cdot \exp(-5)$
Diversity	$4.57 \cdot \exp(-4)$	$1.46 \cdot \exp(-4)$	0.00179 **	-
Age	$1.47 \cdot \exp(-3)$	$2.09 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.84 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-9.60	$7.78 \cdot \exp(1)$	0.902	-
Asia	-	-	-	-
Europe	$-5.04 \cdot \exp(-1)$	$5.98 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.52 \cdot \exp(-6)$
Oceania	$-1.92 \cdot \exp(-1)$	1.00	0.849	-
AIC: 43,894, BIC: 44,098				

TABLE H.6: ERGM Results for Network Snapshot on 16/12/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-9.76	$9.90 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$5.77 \cdot \exp(-5)$
Vendor Degree	$-5.07 \cdot \exp(-4)$	$1.16 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.77 \cdot \exp(-5)$
Trustworthiness	$3.27 \cdot \exp(-3)$	$3.23 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.79 \cdot \exp(-5)$
Earnings	$1.05 \cdot \exp(-4)$	$5.87 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$5.77 \cdot \exp(-5)$
Affordability	$5.58 \cdot \exp(-4)$	$4.34 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.77 \cdot \exp(-5)$
Diversity	$7.83 \cdot \exp(-4)$	$1.40 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.78 \cdot \exp(-5)$
Age	$1.20 \cdot \exp(-3)$	$1.98 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.78 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.02	1.00	0.310	-
Asia	-	-	-	-
Europe	$-6.71 \cdot \exp(-1)$	$5.42 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.95 \cdot \exp(-5)$
Oceania	-9.82	$7.4 \cdot \exp(1)$	0.894	-
AIC: 49,069, BIC: 49,276				

TABLE H.7: ERGM Results for Network Snapshot on 23/12/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.15 \cdot \exp(1)$	$-4.65 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Vendor Degree	$-4.65 \cdot \exp(-4)$	$2.32 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Trustworthiness	$2.49 \cdot \exp(-3)$	$5.13 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Earnings	$1.17 \cdot \exp(-4)$	$1.20 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Affordability	$4.99 \cdot \exp(-4)$	$1.12 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Diversity	$2.21 \cdot \exp(-3)$	$2.97 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Age	$2.46 \cdot \exp(-3)$	$4.07 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.12 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-1.02 \cdot \exp(1)$	$1.44 \cdot \exp(2)$	0.441	-
Asia	-	-	-	-
Europe	$-8.62 \cdot \exp(-1)$	$1.19 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$4.73 \cdot \exp(-6)$
Oceania	$-1.07 \cdot \exp(1)$	$2.04 \cdot \exp(2)$	0.958	-
AIC: 14,562, BIC: 14,768				

TABLE H.8: ERGM Results for Network Snapshot on 30/12/2013

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-9.89	$8.90 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$5.07 \cdot \exp(-5)$
Vendor Degree	$4.89 \cdot \exp(-4)$	$9.07 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$5.07 \cdot \exp(-5)$
Trustworthiness	$1.28 \cdot \exp(-4)$	$2.12 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.07 \cdot \exp(-5)$
Earnings	$1.37 \cdot \exp(-4)$	$5.05 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$5.07 \cdot \exp(-5)$
Affordability	$4.87 \cdot \exp(-4)$	$5.85 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.07 \cdot \exp(-5)$
Diversity	$1.54 \cdot \exp(-3)$	$1.34 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.08 \cdot \exp(-5)$
Age	$2.36 \cdot \exp(-3)$	$1.66 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.08 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.16	1.00	0.246	-
Asia	-	-	-	-
Europe	$-8.25 \cdot \exp(-1)$	$4.77 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.22 \cdot \exp(-5)$
Oceania	$-1.01 \cdot \exp(1)$	$7.27 \cdot \exp(1)$	0.889	-
AIC: 61,912, BIC: 62,122				

TABLE H.9: ERGM Results for Network Snapshot on 06/01/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.01 \cdot \exp(1)$	$7.85 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Vendor Degree	$4.16 \cdot \exp(-4)$	$1.03 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Trustworthiness	$1.76 \cdot \exp(-4)$	$2.10 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Earnings	$1.00 \cdot \exp(-4)$	$2.98 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Affordability	$5.19 \cdot \exp(-4)$	$3.70 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Diversity	$2.01 \cdot \exp(-3)$	$1.17 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-5)$
Age	$2.00 \cdot \exp(-3)$	$1.50 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.19	$7.08 \cdot \exp(-1)$	0.0941 .	-
Asia	$-6.69 \cdot \exp(-1)$	1.00	0.504	-
Europe	$-7.56 \cdot \exp(1)$	$4.18 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.93 \cdot \exp(-5)$
Oceania	-8.99	4.22	0.831	-
AIC: 82,712, BIC: 82,924				

TABLE H.10: ERGM Results for Network Snapshot on 13/01/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-9.94	$7.02 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$4.82 \cdot \exp(-5)$
Vendor Degree	$4.53 \cdot \exp(-4)$	$8.45 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.82 \cdot \exp(-5)$
Trustworthiness	$1.75 \cdot \exp(-4)$	$1.43 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.82 \cdot \exp(-5)$
Earnings	$8.84 \cdot \exp(-5)$	$2.02 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.82 \cdot \exp(-5)$
Affordability	$1.17 \cdot \exp(-4)$	$2.77 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.82 \cdot \exp(-5)$
Diversity	$1.98 \cdot \exp(-3)$	$1.05 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.83 \cdot \exp(-5)$
Age	$1.29 \cdot \exp(-3)$	$1.38 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.83 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	2.03	$1.22 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	0.000367
Asia	-	-	-	-
Europe	$-4.47 \cdot \exp(-1)$	$3.75 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.08 \cdot \exp(-5)$
Oceania	$-4.24 \cdot \exp(-1)$	1.00	0.659	-
AIC: 101,728, BIC: 101,943				

TABLE H.11: ERGM Results for Network Snapshot on 20/01/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.02 \cdot \exp(1)$	$8.39 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.72 \cdot \exp(-5)$
Vendor Degree	$4.55 \cdot \exp(-4)$	$9.26 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$3.72 \cdot \exp(-5)$
Trustworthiness	$1.03 \cdot \exp(-4)$	$1.41 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$3.72 \cdot \exp(-5)$
Earnings	$7.27 \cdot \exp(-5)$	$1.49 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$3.72 \cdot \exp(-5)$
Affordability	$2.03 \cdot \exp(-4)$	$2.87 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$3.72 \cdot \exp(-5)$
Diversity	$2.54 \cdot \exp(-3)$	$1.13 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-5)$
Age	$7.93 \cdot \exp(-4)$	$1.59 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$6.84 \cdot \exp(-1)$	$2.45 \cdot \exp(-1)$	0.00519 **	-
Asia	-	-	-	-
Europe	$8.69 \cdot \exp(-1)$	$4.41 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$8.86 \cdot \exp(-5)$
Oceania	-	-	-	-
AIC: 89,868, BIC: 90,084				

TABLE H.12: ERGM Results for Network Snapshot on 27/01/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.01 \cdot \exp(1)$	$8.04 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Vendor Degree	$4.47 \cdot \exp(-4)$	$9.63 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Trustworthiness	$1.56 \cdot \exp(-4)$	$1.42 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Earnings	$4.64 \cdot \exp(-5)$	$1.10 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Affordability	$3.12 \cdot \exp(-4)$	$2.84 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Diversity	$2.21 \cdot \exp(-3)$	$1.09 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-5)$
Age	$2.38 \cdot \exp(-4)$	$1.56 \cdot \exp(-4)$	0.125	-
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$3.19 \cdot \exp(-1)$	$3.03 \cdot \exp(-1)$	0.294	-
Asia	-	-	-	-
Europe	$-8.38 \cdot \exp(-1)$	$4.19 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.78 \cdot \exp(-5)$
Oceania	$-7.24 \cdot \exp(-1)$	$9.96 \cdot \exp(-1)$	0.467	-
AIC: 99,992, BIC: 100,209				

TABLE H.13: ERGM Results for Network Snapshot on 03/02/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.00 \cdot \exp(1)$	$8.87 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$4.54 \cdot \exp(-5)$
Vendor Degree	$4.43 \cdot \exp(-4)$	$1.08 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.54 \cdot \exp(-5)$
Trustworthiness	$1.36 \cdot \exp(-4)$	$1.41 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.54 \cdot \exp(-5)$
Earnings	$3.42 \cdot \exp(-5)$	$1.18 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.54 \cdot \exp(-5)$
Affordability	$4.56 \cdot \exp(-4)$	$2.91 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.54 \cdot \exp(-5)$
Diversity	$1.54 \cdot \exp(-3)$	$1.17 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.55 \cdot \exp(-5)$
Age	$-1.52 \cdot \exp(-4)$	$1.69 \cdot \exp(-4)$	0.367	-
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$6.48 \cdot \exp(-1)$	$2.79 \cdot \exp(-1)$	0.0204 *	-
Asia	-	-	-	-
Europe	$-7.34 \cdot \exp(-1)$	$4.43 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.18 \cdot \exp(-5)$
Oceania	-9.36	$5.01 \cdot \exp(1)$	0.852	-
AIC: 89,773, BIC: 89,991				

TABLE H.14: ERGM Results for Network Snapshot on 10/02/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.01 \cdot \exp(1)$	$1.46 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Vendor Degree	$4.78 \cdot \exp(-4)$	$1.46 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Trustworthiness	$1.02 \cdot \exp(-4)$	$1.83 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Earnings	$1.55 \cdot \exp(-5)$	$1.82 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Affordability	$2.55 \cdot \exp(-4)$	$4.46 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-5)$
Diversity	$1.85 \cdot \exp(-3)$	$1.82 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-5)$
Age	$-9.70 \cdot \exp(-4)$	$2.70 \cdot \exp(-4)$	0.000334 ***	$4.10 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	1.78	$2.34 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	0.244
Asia	-	-	-	-
Europe	-1.10	$7.15 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Oceania	-	-	-	-
AIC: 45,393, BIC: 45,609				

TABLE H.15: ERGM Results for Network Snapshot on 17/02/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.25 \cdot \exp(1)$	$2.07 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-6)$
Vendor Degree	$4.70 \cdot \exp(-4)$	$1.82 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-6)$
Trustworthiness	$9.90 \cdot \exp(-5)$	$2.36 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-6)$
Earnings	$8.89 \cdot \exp(-6)$	$2.07 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-6)$
Affordability	$2.61 \cdot \exp(-4)$	$5.47 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$3.73 \cdot \exp(-6)$
Diversity	$4.09 \cdot \exp(-3)$	$2.54 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$3.74 \cdot \exp(-6)$
Age	$2.94 \cdot \exp(-3)$	$3.33 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$3.74 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$2.66 \cdot \exp(-1)$	$5.80 \cdot \exp(-1)$	0.646	-
Asia	-	-	-	-
Europe	$8.45 \cdot \exp(-1)$	$8.78 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$8.68 \cdot \exp(-6)$
Oceania	-	-	-	-
AIC: 30,000 BIC: 30,216				

TABLE H.16: ERGM Results for Network Snapshot on 24/02/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.12 \cdot \exp(1)$	$1.47 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Vendor Degree	$3.05 \cdot \exp(-4)$	$2.18 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Trustworthiness	$3.14 \cdot \exp(-4)$	$2.55 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Earnings	$1.06 \cdot \exp(-5)$	$1.50 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Affordability	$2.56 \cdot \exp(-4)$	$5.22 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Diversity	$3.30 \cdot \exp(-3)$	$1.85 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Age	$1.45 \cdot \exp(-3)$	$2.51 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-2.96 \cdot \exp(-1)$	$5.79 \cdot \exp(-1)$	0.609	-
Asia	-	-	-	-
Europe	$-9.74 \cdot \exp(-1)$	$6.61 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$5.16 \cdot \exp(-6)$
Oceania	-1.01	1.00	0.312	-
AIC: 51,406, BIC: 51,623				

TABLE H.17: ERGM Results for Network Snapshot on 03/03/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.10 \cdot \exp(1)$	$1.38 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Vendor Degree	$4.51 \cdot \exp(-4)$	$1.23 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Trustworthiness	$1.23 \cdot \exp(-4)$	$1.41 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Earnings	$1.58 \cdot \exp(-5)$	$1.37 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Affordability	$1.79 \cdot \exp(-4)$	$4.06 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Diversity	$2.32 \cdot \exp(-3)$	$1.79 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Age	$1.26 \cdot \exp(-3)$	$2.33 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.18	1.00	0.240	-
Asia	-	-	-	-
Europe	$-8.00 \cdot \exp(-1)$	$5.99 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$7.50 \cdot \exp(-6)$
Oceania	-	-	-	-
AIC: 54,296, BIC: 54,513				

Table H.18: ERGM Results for Network Snapshot on 10/03/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.18 \cdot \exp(1)$	$1.57 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$7.50 \cdot \exp(-6)$
Vendor Degree	$4.89 \cdot \exp(-4)$	$1.55 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$7.51 \cdot \exp(-6)$
Trustworthiness	$7.86 \cdot \exp(-5)$	$1.67 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$7.51 \cdot \exp(-6)$
Earnings	$1.04 \cdot \exp(-5)$	$1.40 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$7.50 \cdot \exp(-6)$
Affordability	$2.54 \cdot \exp(-4)$	$4.27 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$7.51 \cdot \exp(-6)$
Diversity	$3.31 \cdot \exp(-3)$	$1.99 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$7.53 \cdot \exp(-6)$
Age	$2.36 \cdot \exp(-3)$	$2.54 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$7.52 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.10	1.00	0.270	-
Asia	-	-	-	-
Europe	$-7.64 \cdot \exp(-1)$	$6.38 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.50 \cdot \exp(-6)$
Oceania	$-1.02 \cdot \exp(1)$	$7.19 \cdot \exp(1)$	0.887	-
AIC: 50,357, BIC: 50,574				

Table H.19: ERGM Results for Network Snapshot on 17/03/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.19 \cdot \exp(1)$	$1.54 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$6.79 \cdot \exp(-6)$
Vendor Degree	$1.17 \cdot \exp(-4)$	$3.21 \cdot \exp(-5)$	0.000277 ***	$6.79 \cdot \exp(-6)$
Trustworthiness	$4.63 \cdot \exp(-4)$	$3.19 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.79 \cdot \exp(-6)$
Earnings	$5.33 \cdot \exp(-6)$	$1.36 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$6.79 \cdot \exp(-6)$
Affordability	$3.34 \cdot \exp(-4)$	$4.32 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.79 \cdot \exp(-6)$
Diversity	$3.43 \cdot \exp(-3)$	$1.97 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$6.81 \cdot \exp(-6)$
Age	$2.79 \cdot \exp(-3)$	$2.46 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$6.81 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-5.09 \cdot \exp(-1)$	$7.09 \cdot \exp(-1)$	0.472	-
Asia	-	-	-	-
Europe	$-7.34 \cdot \exp(-1)$	$6.00 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.26 \cdot \exp(-6)$
Oceania	$-1.03 \cdot \exp(1)$	$7.25 \cdot \exp(1)$	0.887	-
AIC: 53,317, BIC: 53,535				

Table H.20: ERGM Results for Network Snapshot on 24/03/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.21 \cdot \exp(1)$	$1.60 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$5.56 \cdot \exp(-6)$
Vendor Degree	$3.31 \cdot \exp(-4)$	$2.26 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.56 \cdot \exp(-6)$
Trustworthiness	$2.43 \cdot \exp(-4)$	$2.16 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.56 \cdot \exp(-6)$
Earnings	$5.89 \cdot \exp(-6)$	$1.34 \cdot \exp(-6)$	$< 1 \cdot \exp(-4)$ ***	$5.56 \cdot \exp(-6)$
Affordability	$3.42 \cdot \exp(-4)$	$4.20 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$5.56 \cdot \exp(-6)$
Diversity	$3.62 \cdot \exp(-3)$	$2.04 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.58 \cdot \exp(-6)$
Age	$2.89 \cdot \exp(-3)$	$2.53 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$5.58 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-4.61 \cdot \exp(-2)$	$5.80 \cdot \exp(-1)$	0.940	-
Asia	-	-	-	-
Europe	$-6.96 \cdot \exp(-1)$	$6.05 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.77 \cdot \exp(-6)$
Oceania	$-1.03 \cdot \exp(1)$	$7.23 \cdot \exp(1)$	0.887	-
AIC: 54,279, BIC: 54,497				

TABLE H.21: ERGM Results for Network Snapshot on 31/03/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.24 \cdot \exp(1)$	$1.63 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Vendor Degree	$3.26 \cdot \exp(-4)$	$1.82 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Trustworthiness	$2.23 \cdot \exp(-4)$	$1.66 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Earnings	$3.43 \cdot \exp(-7)$	$1.27 \cdot \exp(-6)$	0.788	-
Affordability	$4.01 \cdot \exp(-4)$	$3.89 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Diversity	$4.66 \cdot \exp(-3)$	$2.06 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.14 \cdot \exp(-6)$
Age	$3.53 \cdot \exp(-3)$	$2.55 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.13 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-4.47 \cdot \exp(-4)$	$5.80 \cdot \exp(-1)$	0.999	-
Asia	-	-	-	-
Europe	$-8.77 \cdot \exp(-1)$	$6.14 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.71 \cdot \exp(-6)$
Oceania	-1.07	1.00	0.287	-
AIC: 55,824, BIC: 56,042				

TABLE H.22: ERGM Results for Network Snapshot on 07/04/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.24 \cdot \exp(1)$	$1.67 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Vendor Degree	$5.44 \cdot \exp(-4)$	$1.73 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Trustworthiness	$-8.63 \cdot \exp(-6)$	$1.51 \cdot \exp(-5)$	0.566	-
Earnings	$4.02 \cdot \exp(-6)$	$1.25 \cdot \exp(-6)$	0.00133 **	-
Affordability	$4.32 \cdot \exp(-4)$	$4.02 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Diversity	$4.30 \cdot \exp(-3)$	$2.11 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.14 \cdot \exp(-6)$
Age	$3.13 \cdot \exp(-3)$	$2.61 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.13 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	1.07	$3.57 \cdot \exp(-1)$	0.00282 **	-
Asia	-	-	-	-
Europe	$-8.03 \cdot \exp(-1)$	$6.15 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Oceania	$-1.03 \cdot \exp(1)$	$6.30 \cdot \exp(1)$	0.871	-
AIC: 56,899, BIC: 57,118				

TABLE H.23: ERGM Results for Network Snapshot on 14/04/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.30 \cdot \exp(1)$	$1.66 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Vendor Degree	$5.50 \cdot \exp(-4)$	$1.64 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Trustworthiness	$2.04 \cdot \exp(-5)$	$1.39 \cdot \exp(-5)$	0.141	-
Earnings	$-1.31 \cdot \exp(-6)$	$1.10 \cdot \exp(-6)$	0.237	-
Affordability	$5.75 \cdot \exp(-4)$	$4.15 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Diversity	$5.39 \cdot \exp(-3)$	$1.99 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Age	$4.08 \cdot \exp(-3)$	$2.51 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-4.22 \cdot \exp(-1)$	$7.10 \cdot \exp(-1)$	0.552	-
Asia	-	-	-	-
Europe	$-8.33 \cdot \exp(-1)$	$6.06 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$9.83 \cdot \exp(-7)$
Oceania	$-1.04 \cdot \exp(1)$	$6.33 \cdot \exp(1)$	0.870	-
AIC: 61,114, BIC: 61,333				

TABLE H.24: ERGM Results for Network Snapshot on 21/04/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.40 \cdot \exp(1)$	$1.69 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$8.32 \cdot \exp(-7)$
Vendor Degree	$5.77 \cdot \exp(-4)$	$1.86 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$8.32 \cdot \exp(-7)$
Trustworthiness	$-5.85 \cdot \exp(-6)$	$1.52 \cdot \exp(-5)$	0.701	-
Earnings	$-7.47 \cdot \exp(-6)$	$9.37 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$8.32 \cdot \exp(-7)$
Affordability	$7.06 \cdot \exp(-4)$	$4.10 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$8.32 \cdot \exp(-7)$
Diversity	$6.59 \cdot \exp(-3)$	$1.74 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$8.37 \cdot \exp(-7)$
Age	$5.50 \cdot \exp(-3)$	$2.33 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$8.36 \cdot \exp(-7)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$4.05 \cdot \exp(-2)$	$5.80 \cdot \exp(-1)$	0.944	-
Asia	-	-	-	-
Europe	$-7.96 \cdot \exp(-1)$	$6.34 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.57 \cdot \exp(-7)$
Oceania	-	-	-	-
AIC: 56,847, BIC: 57,066				

TABLE H.25: ERGM Results for Network Snapshot on 28/04/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.42 \cdot \exp(1)$	$1.63 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$6.81 \cdot \exp(-7)$
Vendor Degree	$5.11 \cdot \exp(-4)$	$1.97 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.81 \cdot \exp(-7)$
Trustworthiness	$5.29 \cdot \exp(-5)$	$1.53 \cdot \exp(-5)$	0.000522 ***	$6.81 \cdot \exp(-7)$
Earnings	$-1.10 \cdot \exp(-5)$	$8.34 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$6.81 \cdot \exp(-7)$
Affordability	$7.68 \cdot \exp(-4)$	$4.06 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.81 \cdot \exp(-7)$
Diversity	$7.39 \cdot \exp(-3)$	$1.64 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$6.86 \cdot \exp(-7)$
Age	$5.95 \cdot \exp(-3)$	$2.24 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$6.85 \cdot \exp(-7)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-8.22 \cdot \exp(-2)$	$5.80 \cdot \exp(-1)$	0.887	-
Asia	-	-	-	-
Europe	$-9.20 \cdot \exp(-1)$	$6.18 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.71 \cdot \exp(-7)$
Oceania	$-1.04 \cdot \exp(1)$	$6.32 \cdot \exp(1)$	0.869	-
AIC: 63,679, BIC: 63,898				

TABLE H.26: ERGM Results for Network Snapshot on 05/05/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.35 \cdot \exp(1)$	$1.58 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Vendor Degree	$5.97 \cdot \exp(-4)$	$2.04 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Trustworthiness	$-1.22 \cdot \exp(-5)$	$1.51 \cdot \exp(-5)$	0.419	-
Earnings	$-1.12 \cdot \exp(-5)$	$7.82 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Affordability	$6.79 \cdot \exp(-4)$	$4.05 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Diversity	$6.59 \cdot \exp(-3)$	$1.61 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.53 \cdot \exp(-6)$
Age	$5.04 \cdot \exp(-3)$	$2.22 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$5.98 \cdot \exp(-1)$	$4.11 \cdot \exp(-1)$	0.146	-
Asia	-	-	-	-
Europe	-1.04	$6.10 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$5.36 \cdot \exp(-7)$
Oceania	$-1.05 \cdot \exp(1)$	$6.31 \cdot \exp(1)$	0.869	-
AIC: 65,018, BIC: 65,238				

Table H.27: ERGM Results for Network Snapshot on 12/05/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.33 \cdot \exp(1)$	$1.60 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Vendor Degree	$4.02 \cdot \exp(-4)$	$2.20 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Trustworthiness	$1.36 \cdot \exp(-4)$	$1.53 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Earnings	$-5.37 \cdot \exp(-6)$	$7.27 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Affordability	$7.00 \cdot \exp(-4)$ ***	$4.09 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$	$1.68 \cdot \exp(-6)$
Diversity	$5.56 \cdot \exp(-3)$	$1.56 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Age	$4.23 \cdot \exp(-3)$	$2.18 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	1.08	$3.60 \cdot \exp(-1)$	0.00256 **	-
Asia	-	-	-	-
Europe	$-8.35 \cdot \exp(-1)$	$6.19 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$7.27 \cdot \exp(-7)$
Oceania	$-1.03 \cdot \exp(1)$	$6.38 \cdot \exp(1)$	0.872	-
AIC: 61,621, BIC: 61,841				

Table H.28: ERGM Results for Network Snapshot on 19/05/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.34 \cdot \exp(1)$	$1.52 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Vendor Degree	$5.53 \cdot \exp(-4)$	$2.28 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Trustworthiness	$2.53 \cdot \exp(-5)$	$1.50 \cdot \exp(-5)$	0.0902 .	-
Earnings	$-9.54 \cdot \exp(-6)$	$6.79 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Affordability	$7.39 \cdot \exp(-4)$	$3.76 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Diversity	$6.22 \cdot \exp(-3)$	$1.46 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Age	$4.78 \cdot \exp(-3)$	$2.08 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$5.68 \cdot \exp(-1)$	$4.12 \cdot \exp(-1)$	0.168	-
Asia	-	-	-	-
Europe	$-9.82 \cdot \exp(-1)$	$5.84 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$5.68 \cdot \exp(-7)$
Oceania	$-1.05 \cdot \exp(1)$	$6.37 \cdot \exp(1)$	0.869	-
AIC: 69,161, BIC: 69,381				

Table H.29: ERGM Results for Network Snapshot on 26/05/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.33 \cdot \exp(1)$	$1.59 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Vendor Degree	$4.57 \cdot \exp(-4)$	$2.86 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Trustworthiness	$8.96 \cdot \exp(-5)$	$1.77 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Earnings	$-8.81 \cdot \exp(-6)$	$6.94 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Affordability	$8.78 \cdot \exp(-4)$	$4.00 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Diversity	$5.69 \cdot \exp(-3)$	$1.52 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Age	$4.41 \cdot \exp(-3)$	$2.16 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.13	1.00	0.258	-
Asia	-	-	-	-
Europe	$-9.19 \cdot \exp(-1)$	$6.07 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$6.68 \cdot \exp(-7)$
Oceania	-1.26	1.00	0.209	-
AIC: 62,947, BIC: 63,168				

TABLE H.30: ERGM Results for Network Snapshot on 02/06/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.34 \cdot \exp(1)$	$1.56 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Vendor Degree	$4.50 \cdot \exp(-4)$	$2.98 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Trustworthiness	$8.72 \cdot \exp(-5)$	$1.75 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Earnings	$-8.67 \cdot \exp(-6)$	$6.57 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Affordability	$9.12 \cdot \exp(-4)$	$3.90 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Diversity	$5.65 \cdot \exp(-3)$	$1.46 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Age	$4.50 \cdot \exp(-3)$	$2.10 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$1.88 \cdot \exp(-1)$	$4.50 \cdot \exp(-1)$	0.676	-
Asia	-	-	-	-
Europe	$-8.30 \cdot \exp(-1)$	$5.99 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$6.61 \cdot \exp(-7)$
Oceania	$-1.03 \cdot \exp(1)$	$5.70 \cdot exp(1)$	0.856	-
AIC: 65,148, BIC: 65,369				

TABLE H.31: ERGM Results for Network Snapshot on 09/06/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.33 \cdot \exp(1)$	$1.63 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Vendor Degree	$-3.61 \cdot \exp(-4)$	$4.40 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Trustworthiness	$5.48 \cdot \exp(-4)$	$2.47 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Earnings	$-1.01 \cdot \exp(-5)$	$7.14 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Affordability	$9.66 \cdot \exp(-4)$	$4.25 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Diversity	$4.84 \cdot \exp(-3)$	$1.54 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Age	$4.10 \cdot \exp(-3)$	$2.17 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$7.71 \cdot \exp(-1)$	$3.57 \cdot \exp(-1)$	0.0311 *	-
Asia	-	-	-	-
Europe	$-7.31 \cdot \exp(-1)$	$6.12 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$8.06 \cdot \exp(-7)$
Oceania	$-6.05 \cdot \exp(-1)$	$7.09 \cdot \exp(-1)$	0.393	-
AIC: 57,795, BIC: 58,015				

TABLE H.32: ERGM Results for Network Snapshot on 16/06/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.38 \cdot \exp(1)$	$1.61 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Vendor Degree	$3.79 \cdot \exp(-4)$	$3.55 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Trustworthiness	$1.35 \cdot \exp(-4)$	$1.90 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Earnings	$-8.78 \cdot \exp(-6)$	$6.68 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Affordability	$9.59 \cdot \exp(-4)$	$4.23 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Diversity	$5.27 \cdot \exp(-3)$	$1.46 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Age	$4.55 \cdot \exp(-3)$	$2.08 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-6.05 \cdot \exp(-1)$	$7.09 \cdot \exp(-1)$	0.394	-
Asia	-	-	-	-
Europe	$-6.05 \cdot \exp(-1)$	$5.99 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$5.55 \cdot \exp(-7)$
Oceania	-1.46	$9.93 \cdot \exp(-1)$	0.142	-
AIC: 60,509, BIC: 60,730				

TABLE H.33: ERGM Results for Network Snapshot on 23/06/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.34 \cdot \exp(1)$	$1.57 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Vendor Degree	$1.71 \cdot \exp(-4)$	$3.75 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Trustworthiness	$2.30 \cdot \exp(-4)$	$1.91 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Earnings	$-7.40 \cdot \exp(-6)$	$6.42 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Affordability	$9.51 \cdot \exp(-4)$	$4.19 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Diversity	$4.87 \cdot \exp(-3)$	$1.45 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Age	$4.20 \cdot \exp(-3)$	$2.05 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-4.34 \cdot \exp(-2)$	$5.03 \cdot \exp(-1)$	0.931	-
Asia	-	-	-	-
Europe	$-6.80 \cdot \exp(-1)$	$5.79 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$7.68 \cdot \exp(-7)$
Oceania	$-1.03 \cdot \exp(1)$	$5.19 \cdot \exp(1)$	0.843	-
AIC:62,580, BIC: 62,801				

TABLE H.34: ERGM Results for Network Snapshot on 30/06/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.33 \cdot \exp(1)$	$1.72 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Vendor Degree	$-6.85 \cdot \exp(-4)$	$3.51 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Trustworthiness	$6.33 \cdot \exp(-4)$	$1.68 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Earnings	$9.83 \cdot \exp(-6)$	$9.10 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.67 \cdot \exp(-6)$
Affordability	$6.59 \cdot \exp(-4)$	$5.70 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Diversity	$4.23 \cdot \exp(-3)$	$1.53 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Age	$3.50 \cdot \exp(-3)$	$2.22 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.68 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$7.80 \cdot \exp(-1)$	$3.82 \cdot \exp(-1)$	0.0409 .	-
Asia	-	-	-	-
Europe	$-5.86 \cdot \exp(-1)$	$6.37 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$9.32 \cdot \exp(-7)$
Oceania	-9.97	$5.22 \cdot \exp(1)$	0.849	-
AIC: 53,860, BIC: 54,080				

TABLE H.35: ERGM Results for Network Snapshot on 07/07/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.32 \cdot \exp(1)$	$1.63 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Vendor Degree	$-4.80 \cdot \exp(-4)$	$3.63 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Trustworthiness	$5.16 \cdot \exp(-4)$	$1.67 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Earnings	$6.34 \cdot \exp(-6)$	$7.76 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Affordability	$7.13 \cdot \exp(-4)$	$5.44 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Diversity	$4.49 \cdot \exp(-3)$	$1.38 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.86 \cdot \exp(-6)$
Age	$3.46 \cdot \exp(-3)$	$2.10 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.86 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-4.28 \cdot \exp(-2)$	$5.80 \cdot \exp(-1)$	0.941	-
Asia	-	-	-	-
Europe	$-7.17 \cdot \exp(-1)$	$6.20 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$9.03 \cdot \exp(-7)$
Oceania	$-1.01 \cdot \exp(1)$	$5.20 \cdot \exp(1)$	0.846	-
AIC: 59,279, BIC: 59,499				

TABLE H.36: ERGM Results for Network Snapshot on 14/07/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.32 \cdot \exp(1)$	$1.68 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Vendor Degree	$-4.89 \cdot \exp(-4)$	$3.60 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Trustworthiness	$4.89 \cdot \exp(-4)$	$1.56 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Earnings	$7.15 \cdot \exp(-6)$	$7.48 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Affordability	$7.61 \cdot \exp(-4)$	$5.05 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Diversity	$4.15 \cdot \exp(-3)$	$1.46 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.86 \cdot \exp(-6)$
Age	$3.29 \cdot \exp(-3)$	$2.13 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.86 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$7.41 \cdot \exp(-1)$	$4.12 \cdot \exp(-1)$	0.0721 .	-
Asia	-	-	-	-
Europe	$-5.55 \cdot \exp(-1)$	$6.18 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$9.46 \cdot \exp(-7)$
Oceania	-	-	-	-
AIC: 59,230, BIC: 59,450				

TABLE H.37: ERGM Results for Network Snapshot on 21/07/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.32 \cdot \exp(1)$	$1.59 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Vendor Degree	$1.74 \cdot \exp(-4)$	$3.39 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Trustworthiness	$1.94 \cdot \exp(-4)$	$1.44 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Earnings	$-1.69 \cdot \exp(-6)$	$6.17 \cdot \exp(-7)$	0.00634 **	-
Affordability	$9.45 \cdot \exp(-4)$	$4.21 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.85 \cdot \exp(-6)$
Diversity	$4.53 \cdot \exp(-3)$	$1.30 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.86 \cdot \exp(-6)$
Age	$3.57 \cdot \exp(-3)$	$2.01 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.86 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$4.23 \cdot \exp(-1)$	$4.51 \cdot \exp(-1)$	0.348	-
Asia	-	-	-	-
Europe	$-6.71 \cdot \exp(-1)$	$6.00 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$9.46 \cdot \exp(-7)$
Oceania	$-1.02 \cdot \exp(1)$	$5.29 \cdot \exp(1)$	0.847	-
AIC: 63,281, BIC: 63,501				

TABLE H.38: ERGM Results for Network Snapshot on 28/07/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.31 \cdot \exp(1)$	$1.51 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Vendor Degree	$5.12 \cdot \exp(-4)$	$2.60 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Trustworthiness	$5.67 \cdot \exp(-5)$	$1.07 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Earnings	$-4.12 \cdot \exp(-6)$	$5.20 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Affordability	$8.64 \cdot \exp(-4)$	$3.80 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Diversity	$4.44 \cdot \exp(-3)$	$1.23 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Age	$3.43 \cdot \exp(-3)$	$1.89 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$2.18 \cdot \exp(-1)$	$4.51 \cdot \exp(-1)$	0.629	-
Asia	-	-	-	-
Europe	$-6.37 \cdot \exp(-1)$	$5.65 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.08 \cdot \exp(-6)$
Oceania	$-1.03 \cdot \exp(1)$	$5.26 \cdot \exp(1)$	0.845	-
AIC: 72,171, BIC: 72,391				

Table H.39: ERGM Results for Network Snapshot on 04/08/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.34 \cdot \exp(1)$	$1.55 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Vendor Degree	$8.14 \cdot \exp(-4)$	$3.23 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Trustworthiness	$-7.13 \cdot \exp(-5)$	$1.27 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Earnings	$-5.65 \cdot \exp(-6)$	$5.17 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Affordability	$8.67 \cdot \exp(-4)$	$3.69 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Diversity	$4.38 \cdot \exp(-3)$	$1.26 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Age	$3.74 \cdot \exp(-3)$	$1.89 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.52 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-5.40 \cdot \exp(-1)$	$7.10 \cdot \exp(-1)$	0.447	-
Asia	-	-	-	-
Europe	$-4.59 \cdot \exp(-1)$	$5.74 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$9.57 \cdot \exp(-7)$
Oceania	$-1.02 \cdot \exp(1)$	$5.24 \cdot \exp(1)$	0.846	-
AIC: 68,427, BIC: 68,647				

Table H.40: ERGM Results for Network Snapshot on 11/08/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.30 \cdot \exp(1)$	$1.60 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Vendor Degree	$5.86 \cdot \exp(-4)$	$3.37 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Trustworthiness	$1.21 \cdot \exp(-5)$	$1.24 \cdot \exp(-5)$	0.330	-
Earnings	$-4.70 \cdot \exp(-6)$	$4.74 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Affordability	$1.02 \cdot \exp(-3)$	$3.92 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Diversity	$4.25 \cdot \exp(-3)$	$1.22 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Age	$3.28 \cdot \exp(-3)$	$1.96 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-5.87 \cdot \exp(-1)$	$7.09 \cdot \exp(-1)$	0.408	-
Asia	-	-	-	-
Europe	$-6.31 \cdot \exp(-1)$	$6.14 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.20 \cdot \exp(-6)$
Oceania	-1.40	$9.94 \cdot \exp(-1)$	0.159	-
AIC: 64,932, BIC: 65,152				

Table H.41: ERGM Results for Network Snapshot on 18/08/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.30 \cdot \exp(1)$	$1.59 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Vendor Degree	$5.16 \cdot \exp(-4)$	$3.37 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Trustworthiness	$3.24 \cdot \exp(-5)$	$1.19 \cdot \exp(-5)$	0.00624 **	-
Earnings	$-4.42 \cdot \exp(-6)$	$4.41 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Affordability	$1.03 \cdot \exp(-3)$	$3.89 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Diversity	$4.07 \cdot \exp(-3)$	$1.19 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Age	$3.35 \cdot \exp(-3)$	$1.91 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$4.29 \cdot \exp(-1)$	$3.81 \cdot \exp(-1)$	0.260	-
Asia	-	-	-	-
Europe	$-5.77 \cdot \exp(-1)$	$6.08 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.27 \cdot \exp(-6)$
Oceania	$-1.02 \cdot \exp(1)$	$4.91 \cdot \exp(-1)$	0.836	-
AIC: 63,927, BIC: 64,147				

TABLE H.42: ERGM Results for Network Snapshot on 25/08/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.29 \cdot \exp(1)$	$1.54 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Vendor Degree	$5.97 \cdot \exp(-4)$	$3.08 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Trustworthiness	$5.92 \cdot \exp(-6)$	$1.06 \cdot \exp(-5)$	0.577	-
Earnings	$-4.55 \cdot \exp(-6)$	$4.02 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Affordability	$1.04 \cdot \exp(-3)$	$3.70 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Diversity	$3.94 \cdot \exp(-3)$	$1.14 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.51 \cdot \exp(-6)$
Age	$3.30 \cdot \exp(-3)$	$1.84 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.51 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$2.93 \cdot \exp(-1)$	$4.11 \cdot \exp(-1)$	0.476	-
Asia	-	-	-	-
Europe	$-5.63 \cdot \exp(-1)$	$5.89 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.42 \cdot \exp(-6)$
Oceania	$-1.03 \cdot \exp(1)$	$4.89 \cdot \exp(1)$	0.834	-
AIC: 67,611, BIC: 67,830				

TABLE H.43: ERGM Results for Network Snapshot on 01/09/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.31 \cdot \exp(1)$	$1.51 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Vendor Degree	$7.06 \cdot \exp(-4)$	$3.43 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Trustworthiness	$-3.25 \cdot \exp(-5)$	$1.17 \cdot \exp(-5)$	0.00544 **	-
Earnings	$-5.17 \cdot \exp(-6)$	$3.66 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Affordability	$1.04 \cdot \exp(-3)$	$3.80 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Diversity	$3.92 \cdot \exp(-3)$	$1.08 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Age	$3.75 \cdot \exp(-3)$	$1.76 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.05 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$3.44 \cdot \exp(-1)$	$5.04 \cdot \exp(-1)$	0.495	-
Asia	-	-	-	-
Europe	$-4.61 \cdot \exp(-1)$	$5.79 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.29 \cdot \exp(-6)$
Oceania	$-5.09 \cdot \exp(-1)$	$5.79 \cdot \exp(-1)$	0.380	-
AIC: 65,107, BIC: 65,326				

TABLE H.44: ERGM Results for Network Snapshot on 08/09/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.38 \cdot \exp(1)$	$1.75 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Vendor Degree	$4.24 \cdot \exp(-3)$	$1.01 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Trustworthiness	$-1.23 \cdot \exp(-3)$	$3.47 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.01 \cdot \exp(-6)$
Earnings	$-1.69 \cdot \exp(-5)$	$4.88 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Affordability	$1.02 \cdot \exp(-3)$	$4.96 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Diversity	$4.66 \cdot \exp(-3)$	$1.20 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Age	$4.63 \cdot \exp(-3)$	$2.04 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.02 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-1.68 \cdot \exp(-1)$	1.00	0.867	-
Asia	-	-	-	-
Europe	$-5.58 \cdot \exp(-1)$	$6.74 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$5.81 \cdot \exp(-7)$
Oceania	-1.44	$9.91 \cdot \exp(-1)$	0.148	-
AIC: 48,410, BIC: 48,628				

TABLE H.45: ERGM Results for Network Snapshot on 15/09/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.35 \cdot \exp(1)$	$1.40 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-6)$
Vendor Degree	$4.79 \cdot \exp(-3)$	$8.80 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.38 \cdot \exp(-6)$
Trustworthiness	$-1.34 \cdot \exp(-3)$	$2.97 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-6)$
Earnings	$-1.66 \cdot \exp(-5)$	$3.67 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-6)$
Affordability	$9.91 \cdot \exp(-4)$	$4.25 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.37 \cdot \exp(-6)$
Diversity	$4.77 \cdot \exp(-3)$	$9.37 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$1.38 \cdot \exp(-6)$
Age	$4.55 \cdot \exp(-3)$	$1.63 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$1.38 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$6.51 \cdot \exp(-1)$	$4.52 \cdot \exp(-1)$	0.150	-
Asia	-	-	-	-
Europe	$-5.58 \cdot \exp(-1)$	$5.45 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$7.85 \cdot \exp(-7)$
Oceania	$-1.04 \cdot \exp(1)$	$4.38 \cdot \exp(1)$	0.812	-
AIC: 72,802, BIC: 73,020				

TABLE H.46: ERGM Results for Network Snapshot on 22/09/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.30 \cdot \exp(1)$	$1.45 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Vendor Degree	$3.71 \cdot \exp(-3)$	$9.64 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Trustworthiness	$-9.98 \cdot \exp(-4)$	$3.13 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Earnings	$-1.15 \cdot \exp(-5)$	$3.54 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Affordability	$9.76 \cdot \exp(-4)$	$4.14 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.26 \cdot \exp(-6)$
Diversity	$4.20 \cdot \exp(-3)$	$9.73 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Age	$3.83 \cdot \exp(-3)$	$1.66 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.27 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$1.28 \cdot \exp(-1)$	$7.11 \cdot \exp(-1)$	0.858	-
Asia	-	-	-	-
Europe	$-4.96 \cdot \exp(-1)$	$5.63 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.38 \cdot \exp(-6)$
Oceania	$-9.90 \cdot \exp(-2)$	$4.10 \cdot \exp(-1)$	0.809	-
AIC: 73,055, BIC: 73,272				

TABLE H.47: ERGM Results for Network Snapshot on 29/09/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.29 \cdot \exp(1)$	$1.38 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Vendor Degree	$3.14 \cdot \exp(-3)$	$9.07 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Trustworthiness	$-7.87 \cdot \exp(-4)$	$2.85 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Earnings	$-9.81 \cdot \exp(-6)$	$2.96 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Affordability	$9.79 \cdot \exp(-4)$	$3.94 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.50 \cdot \exp(-6)$
Diversity	$3.98 \cdot \exp(-3)$	$9.13 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.51 \cdot \exp(-6)$
Age	$3.85 \cdot \exp(-3)$	$1.56 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.51 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-9.70	$5.47 \cdot \exp(1)$	0.859	-
Asia	-	-	-	-
Europe	$-4.20 \cdot \exp(-1)$	$5.38 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$1.66 \cdot \exp(-6)$
Oceania	$-3.26 \cdot \exp(-1)$	$3.80 \cdot \exp(-1)$	0.391	-
AIC: 76,169, BIC: 76,385				

Appendix I

ERGM Results

Table I.1: ERGM Results for Network Snapshot on 06/10/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.23 \cdot \exp(1)$	$1.42 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$4.55 \cdot \exp(-6)$
Vendor Degree	$2.03 \cdot \exp(-3)$	$9.86 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.56 \cdot \exp(-6)$
Trustworthiness	$-4.37 \cdot \exp(-4)$	$3.00 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.55 \cdot \exp(-6)$
Earnings	$-5.37 \cdot \exp(-6)$	$2.71 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$4.55 \cdot \exp(-6)$
Affordability	$6.49 \cdot \exp(-4)$	$3.48 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.55 \cdot \exp(-6)$
Diversity	$3.52 \cdot \exp(-3)$	$9.06 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.57 \cdot \exp(-6)$
Age	$3.18 \cdot \exp(-3)$	$1.62 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.57 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-7.76 \cdot \exp(-1)$	$9.97 \cdot \exp(-1)$	0.441	-
Asia	-	-	-	-
Europe	$-5.22 \cdot \exp(-1)$	$5.62 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.70 \cdot \exp(-6)$
Oceania	$-5.76 \cdot \exp(-1)$	$4.47 \cdot \exp(-1)$	0.198	-
AIC: 75,649, BIC: 75,864				

Table I.2: ERGM Results for Network Snapshot on 13/10/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.24 \cdot \exp(1)$	$2.34 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Vendor Degree	$4.66 \cdot \exp(-3)$	$1.27 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.14 \cdot \exp(-6)$
Trustworthiness	$-1.19 \cdot \exp(-3)$	$3.77 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.11 \cdot \exp(-6)$
Earnings	$-6.87 \cdot \exp(-6)$	$2.46 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Affordability	$6.08 \cdot \exp(-4)$	$3.58 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.12 \cdot \exp(-6)$
Diversity	$3.80 \cdot \exp(-3)$	$8.40 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$4.13 \cdot \exp(-6)$
Age	$3.85 \cdot \exp(-3)$	$1.50 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$4.13 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$2.21 \cdot \exp(-1)$	$4.12 \cdot \exp(-1)$	0.592	-
Asia	-	-	-	-
Europe	$-4.96 \cdot \exp(-1)$	$5.29 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$2.51 \cdot \exp(-6)$
Oceania	-1.17	$5.78 \cdot \exp(-1)$	0.0438 *	-
AIC: 75,436, BIC: 75,648				

Table I.3: ERGM Results for Network Snapshot on 20/10/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.20 \cdot \exp(1)$	$1.47 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$6.14 \cdot \exp(-6)$
Vendor Degree	$4.45 \cdot \exp(-3)$	$2.55 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$6.17 \cdot \exp(-6)$
Trustworthiness	$-1.11 \cdot \exp(-3)$	$7.32 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.14 \cdot \exp(-6)$
Earnings	$-4.60 \cdot \exp(-6)$	$2.71 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$6.14 \cdot \exp(-6)$
Affordability	$9.01 \cdot \exp(-4)$	$3.89 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.15 \cdot \exp(-6)$
Diversity	$3.31 \cdot \exp(-3)$	$9.99 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$6.16 \cdot \exp(-6)$
Age	$3.09 \cdot \exp(-3)$	$1.68 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$6.16 \cdot \exp(-6)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-3.04 \cdot \exp(-1)$	$4.50 \cdot \exp(-1)$	0.498	-
Asia	-	-	-	-
Europe	$-5.54 \cdot \exp(-1)$	$5.66 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$3.53 \cdot \exp(-6)$
Oceania	$-8.36 \cdot \exp(-1)$	$5.02 \cdot \exp(-1)$	0.0955 .	-
AIC: 72,951, BIC: 73,160				

Table I.4: ERGM Results for Network Snapshot on 27/10/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	$-1.05 \cdot \exp(1)$	$1.4 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$2.75 \cdot \exp(-5)$
Vendor Degree	$-1.32 \cdot \exp(-3)$	$2.44 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.75 \cdot \exp(-5)$
Trustworthiness	$5.29 \cdot \exp(-4)$	$6.84 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.76 \cdot \exp(-5)$
Earnings	$-2.63 \cdot \exp(-6)$	$2.47 \cdot \exp(-7)$	$< 1 \cdot \exp(-4)$ ***	$2.75 \cdot \exp(-5)$
Affordability	$9.25 \cdot \exp(-4)$	$3.49 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$2.76 \cdot \exp(-5)$
Diversity	$2.06 \cdot \exp(-3)$	$9.72 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.76 \cdot \exp(-5)$
Age	$1.92 \cdot \exp(-3)$	$1.57 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$2.76 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	-1.36	1.00	0.175	-
Asia	-	-	-	-
Europe	$-4.33 \cdot \exp(-1)$	$5.36 \cdot \exp(-2)$	$< 1 \cdot \exp(-4)$ ***	$1.79 \cdot \exp(-5)$
Oceania	$-5.03 \cdot \exp(-1)$	$3.81 \cdot \exp(-1)$	0.187	-
AIC: 74,935, BIC: 75,138				

Table I.5: ERGM Results for Network Snapshot on 03/11/2014

Variable	Coefficient	Standard Error	p-value	Cumulative Probability of Tie
Edges	-9.55	$1.72 \cdot \exp(-1)$	$< 1 \cdot \exp(-4)$ ***	$7.12 \cdot \exp(-5)$
Vendor Degree	$-4.66 \cdot \exp(-3)$	$1.69 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$7.09 \cdot \exp(-5)$
Trustworthiness	$1.43 \cdot \exp(-3)$	$4.64 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$7.13 \cdot \exp(-5)$
Earnings	$-8.72 \cdot \exp(-7)$	$2.98 \cdot \exp(-7)$	0.00349 **	-
Affordability	$7.12 \cdot \exp(-4)$	$4.98 \cdot \exp(-5)$	$< 1 \cdot \exp(-4)$ ***	$7.12 \cdot \exp(-5)$
Diversity	$1.06 \cdot \exp(-3)$	$1.15 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$7.13 \cdot \exp(-5)$
Age	$1.43 \cdot \exp(-3)$	$1.89 \cdot \exp(-4)$	$< 1 \cdot \exp(-4)$ ***	$7.13 \cdot \exp(-5)$
Location				
UK	-	-	-	-
Africa	-	-	-	-
America	$-9.96 \cdot \exp(-1)$	$7.09 \cdot \exp(-1)$	0.160	-
Asia	-	-	-	-
Europe	$-2.41 \cdot \exp(-1)$	$6.68 \cdot \exp(-2)$	0.000314 ***	$5.59 \cdot \exp(-5)$
Oceania	$-4.04 \cdot \exp(-1)$	$5.80 \cdot \exp(-1)$	0.486	-
AIC: 41,777, BIC: 41,964				

Bibliography

Afroz, S., Garg, V., McCoy, D., and Greenstadt, R. (2013). Honor among thieves: A common's analysis of cybercrime economies. In *eCrime Researchers Summit (eCRS), 2013*, pages 1–11. IEEE.

Al-Imam, A. (2017). Retrospective analyses of high-risk NPS: Integrative analyses of pubmed, drug fora, and the surface web. *Global Journal of Health Science*, 9(11):40.

Al-Imam, A. and AbdulMajeed, B. A. (2017a). Captagon, octodrine, and NBOMe: An integrative analysis of trends databases, the deep web, and the darknet. *Global Journal of Health Science*, 9(11):114.

Al-Imam, A. and AbdulMajeed, B. A. (2017b). The NPS phenomenon and the deep web: Trends analyses and internet snapshots. *Global Journal of Health Science*, 9(11):71.

Aldridge, J. and Décary-Hétu, D. (2014). Not an 'Ebay for drugs': the cryptomarket 'Silk Road' as a paradigm shifting criminal innovation.

Aldridge, J. and Décary-Hétu, D. (2015). A response to Dolliver's "evaluating drug trafficking on the tor network: Silk road 2, the sequel". *International journal on drug policy*, 26(11):1124–1125.

Ayling, J. (2009). Criminal organizations and resilience. *International Journal of Law, Crime and Justice*, 37(4):182–196.

Bancroft, A. and Reid, P. S. (2016). Concepts of illicit drug quality among darknet market users: Purity, embodied experience, craft and chemical knowledge. *International Journal of Drug Policy*, 35:42–49.

Barabási, A.-L. et al. (2016). *Network science*. Cambridge university press.

Barratt, M. J., Ferris, J. A., and Winstock, A. R. (2014). Use of Silk Road, the online drug marketplace, in the United Kingdom, Australia and the United States. *Addiction*, 109(5):774–783.

Barratt, M. J., Ferris, J. A., and Winstock, A. R. (2016). Safer scoring? cryptomarkets, social supply and drug market violence. *International Journal of Drug Policy*, 35:24–31.

BBC (2014a). Darknet: Bitcoin and drugs worth 1.5m seized by irish police. *bbc. co. uk*.

BBC (2014b). Five arrested in Utopia dark net marketplace crackdown. *bbc. co. uk*.

BBC (2015). Silk Road linked to six drug overdose deaths. *bbc. co. uk*.

BBC (2017). Dark web markets boom after Alphabay and Hansa busts. *bbc. co. uk*.

Bearman, J. and Hanuka, T. (2015a). The untold story of Silk Road, part 1. *wired. com*.

Bearman, J. and Hanuka, T. (2015b). The untold story of Silk Road, part 2: The fall. *wired. com*.

Bhaskar, V., Linacre, R., Machin, S., et al. (2017). The economic functioning of online drugs markets. Technical report, Centre for Economic Performance, LSE.

Biryukov, A. and Pustogarov, I. (2015). Bitcoin over Tor isn't a good idea. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 122–134. IEEE.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.

Blockchain Luxembourg S.A. (2017). Bitcoins in circulation. *blockchain. info*.

Blond, S. L., Manils, P., Abdelberi, C., Kaafar, M. A. D., Castelluccia, C., Legout, A., and Dabbous, W. (2011). One bad apple spoils the bunch: exploiting P2P applications to trace and profile Tor users. *arXiv preprint arXiv:1103.1518*.

Bolger, K. E. and Patterson, C. J. (2003). Sequelae of child maltreatment: Vulnerability and resilience. *Resilience and vulnerability: Adaptation in the context of childhood adversities*, pages 156–181.

Branwen, G. (2013). Darknet market mortality risks. *gwern. net*.

Branwen, G. (2017). Tor dnm-related arrests. *gwern. com*.

Branwen, G., Christin, N., Décary-Hétu, D., Andersen, R., Presidente, E., et al. (2015). Dark net market archives, 2011–2015.

Broséus, J., Rhumorbarbe, D., Mireault, C., Ouellette, V., Crispino, F., and Décary-Hétu, D. (2016). Studying illicit drug trafficking on darknet markets: structure and organisation from a Canadian perspective. *Forensic science international*, 264:7–14.

Broséus, J., Rhumorbarbe, D., Morelato, M., Staehli, L., and Rossy, Q. (2017). A geographical analysis of trafficking on a popular darknet market. *Forensic Science International*.

Buntinx, J. (2016). Operation Hyperion responsible for shutting down several Tor markets. *livebitcoinnews.com*.

Buterin, V. (2013). Freedom hosting taken down, founder arrested, users fed Javascript exploits. *bitcoinhosting.com*.

Buxton, J. and Bingham, T. (2015). The rise and challenge of dark net drug markets. *Policy Brief*, 7.

Caines, A., Pastrana, S., Hutchings, A., and Buttery, P. J. (2018). Automatically identifying the function and intent of posts in underground forums. *Crime Science*, 7:19.

Carver, C. S. (1998). Resilience and thriving: Issues, models, and linkages. *Journal of social issues*, 54(2):245–266.

Caudevilla, F. (2016). The emergence of deep web marketplaces: A health perspective. In *The internet and drug markets (EMCDDA Insights 21)*, pages 69–75. Publications Office of the European Union, Luxembourg.

Chen, A. (2011). The underground website where you can buy any drug. *gawker.com*.

Christin, N. (2013). Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22nd international conference on World Wide Web*, pages 213–224. ACM.

Chu, B., Holt, T. J., and Ahn, G. J. (2010). Examining the creation, distribution, and function of malware on-line. *National Institute of Justice, Washington, DC*.

Chung, F. (2015). 'I think they will start breaking toes': Evolution marketplace goes down in massive bitcoin heist. *news.com.au*.

Ciaian, P., Rajcaniova, M., and Kancs, d. (2016). The economics of Bitcoin price formation. *Applied Economics*, 48(19):1799–1815.

Clark, L. (2013). A guide to the Silk Road shutdown. *wired.com*.

C.M. (2018). Analysis: The effect of dream market's fentanyl ban. *darkwebnews.com*.

Coindesk (2014). Bitcoin-over-Tor anonymity 'can be busted for $2,500 a month'. *reddit.com*.

Coindesk (2017). Bitcoin (USD) price. *coindesk.com*.

Coin.Market (2016a). BTC/USD – Litecoin / US Dollar alltime charts and orderbook. *cryptocoincharts.info*.

Coin.Market (2016b). LTC/USD – Litecoin / US Dollar alltime charts and orderbook. *cryptocoincharts.info*.

CoinMarketCap (2016). Dogecoin (DOGE). *coinmarketcap.com*.

Coopman, V., Cordonnier, J., De Leeuw, M., and Cirimele, V. (2016). Ocfentanil overdose fatality in the recreational drug scene. *Forensic Science International*, 266:469–473.

Corbin, J. M. and Strauss, A. (1990). Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology*, 13(1):3–21.

Côté, I. M. and Darling, E. S. (2010). Rethinking ecosystem resilience in the face of climate change. *PLoS biology*, 8(7):e1000438.

Cox, J. (2016a). Operation Hyperion targets suspected dark web users around the world. *vice.com*.

Cox, J. (2016b). Reputation is everything: the role of ratings, feedback and reviews in cryptomarkets. *The Internet and drug markets*, pages 49–54.

Cumming, G. S., Barnes, G., Perz, S., Schmink, M., Sieving, K. E., Southworth, J., Binford, M., Holt, R. D., Stickler, C., and van Holt, T. (2005). An exploratory framework for the empirical measurement of resilience. *Ecosystems*, 8(8):975–987.

DarkNetMarkets.org (2013). Dark net market information. *darknetmarkets.org*.

DEA (2012). Operation Adam Bomb: Arrest of creators, operators of online secret narcotics marketplace. *DEA.gov*.

Décary-Hétu, D. and Giommoni, L. (2016). Do police crackdowns disrupt drug cryptomarkets? a longitudinal analysis of the effects of Operation Onymous. *Crime, Law and Social Change*, pages 1–21.

DeepDotWeb (2013). Silk Road admin Libertas Inigo arrested. *deepdotweb.com*.

DeepDotWeb (2014a). Another marketplace doxxed: "Hansamarket". *deepdotweb.com*.

DeepDotWeb (2014b). The Majestic Garden. *deepdotweb.com*.

DeepDotWeb (2014c). Pandora marketplace hacked: Losing $250,000 in BTC. *deepdotweb.com*.

DeepDotWeb (2014d). Silk Road 2 hacked, all Bitcoins stolen $2.7 miliion. *deepdotweb.com*.

DeepDotWeb (2014e). Silk Road 2 hacked Bitcoins stolen unknown amount. *deepdotweb.com*.

DeepDotWeb (2014f). Utopia marketplace seized by Dutch police - 5 arrested. *deepdotweb.com*.

DeepDotWeb (2015a). Agora market to stop listing lethal weapons. *deepdotweb.com*.

DeepDotWeb (2015b). Cyruserv hosting shuts down. *deepdotweb.com*.

DeepDotWeb (2015c). East India Market hacked 30 BTC gone. *deepdotweb.com*.

DeepDotWeb (2015d). Evolution marketplace exit scam: Biggest exist scam ever? *deepdotweb.com*.

DeepDotWeb (2015e). Shiny Flakes bust: 38 houses raided. *deepdotweb.com*.

DeepDotWeb (2016a). 10 arrested as part of Operation Hyperion in Sweden. *deepdotweb.com*.

DeepDotWeb (2016b). Austrian police bust a 19-year-old "Powerseller" darknet vender. *deepdotweb.com*.

DeepDotWeb (2016c). German police arrests user who bought weed on Silk Road years ago. *deepdotweb.com*.

DeepDotWeb (2016d). Suspect claims someone else used his computer to buy drugs, jury agrees. *deepdotweb.com*.

DeepDotWeb (2016e). Upper Austria vendor group arrested. *deepdotweb.com*.

Digital Citizens Alliance (2013). Busted, but not broken, the state of Silk Road and the darknet marketplaces. *digitalcitizensalliance.org*.

Digital Citizens Alliance (2014). Busted, but not broken the state of Silk Road and the darknet marketplaces. *A DIGITAL CITIZENS ALLIANCE INVESTIGATIVE REPORT*.

Dittus, M., Wright, J., and Graham, M. (2017). Platform criminalism: The 'last-mile' geography of the darknet market supply chain. *arXiv preprint arXiv:1712.10068*.

Dolliver, D. S. (2015a). Evaluating drug trafficking on the Tor network: Silk Road 2, the sequel. *International Journal of Drug Policy*, 26(11):1113–1123.

Dolliver, D. S. (2015b). A rejoinder to authors: Data collection on Tor. *International Journal of Drug Policy*, 26(11):1128–1129.

Dolliver, D. S. and Kenney, J. L. (2016). Characteristics of drug vendors on the Tor network: a cryptomarket comparison. *Victims & Offenders*, 11(4):600–620.

Drugs Forum User 5-HT2A (2016). 'Operation Hyperion' targets suspected dark web users around the world. *drugs-forum.com*.

DrugWise (2016). What are the UK drug laws? *drugwise.org.uk*.

Durrett, G., Kummerfeld, J. K., Berg-Kirkpatrick, T., Portnoff, R. S., Afroz, S., McCoy, D., Levchenko, K., and Paxson, V. (2017). Identifying products in online cybercrime marketplaces: A dataset for fine-grained domain adaptation. *arXiv preprint arXiv:1708.09609*.

Duxbury, S. W. and Haynie, D. L. (2017). The network structure of opioid distribution on a darknet Cryptomarket. *Journal of Quantitative Criminology*, pages 1–21.

Elmqvist, T., Folke, C., Nyström, M., Peterson, G., Bengtsson, J., Walker, B., and Norberg, J. (2003). Response diversity, ecosystem change, and resilience. *Frontiers in Ecology and the Environment*, 1(9):488–494.

Enlighten (2017). Nintendo. *pillreports.net*.

European Commission (2019). Project to prevent criminal use of the dark web and virtual currencies launched by international consortium. *Cordis*.

Europol (2014). Global action against dark markets on Tor network. *europol.europa.eu*.

Europol (2017). Massive blow to criminal dark web activities after globally coordinated operation. *europol.europa.eu*.

FBI (2016). A primer on darknet marketplaces. *FBI.gov*.

Folke, C. (2006). Resilience: The emergence of a perspective for social–ecological systems analyses. *Global environmental change*, 16(3):253–267.

Franklin, J., Perrig, A., Paxson, V., and Savage, S. (2007). An inquiry into the nature and causes of the wealth of internet miscreants. In *ACM conference on Computer and communications security*, pages 375–388.

Fusion Media Limited (2016a). AUD/USD – Australian Dollar US Dollar. *uk.investing.com*.

Fusion Media Limited (2016b). CAD/USD – Canadian Dollar US Dollar. *uk. investing. com*.

Fusion Media Limited (2016c). EUR/USD – Euro US Dollar. *uk. investing. com*.

Garmezy, N. (1987). Stress, competence, and development: Continuities in the study of schizophrenic adults, children vulnerable to psychopathology, and the search for stress-resistant children. *American journal of Orthopsychiatry*, 57(2):159.

Godschalk, D. R. (2003). Urban hazard mitigation: creating resilient cities. *Natural hazards review*, 4(3):136–143.

Graczyk, M. and Kinningham, K. (2015). Automatic product categorization for anonymous marketplaces.

Greenberg, A. (2013a). Follow the Bitcoins: How we got busted buying drugs on Silk Roads black market. *forbes. com*.

Greenberg, A. (2013b). An interiew with a digital drug lord: the Silk Road's Dread Pirate Roberts. *forbes. com*.

Greenberg, A. (2013c). Silk Road 2.0 launches promising. *forbes. com*.

Greenberg, A. (2013d). Silk Road competitor shuts down and another plans to go offline after $6 million theft. *forbes. com*.

Greenberg, A. (2014a). Five men arrested in Dutch crackdown on Silk Road copycat. *forbes. com*.

Greenberg, A. (2014b). Global web crackdown arrests 17, seizes hundreds of darknet domains. *wired. com*.

Greenberg, A. (2014c). Not just Silk Road 2: Feds seize two other drug markets and counting. *wired. com*.

Greenberg, A. (2014d). Schumer crackdown on dark web drug sales. *wired. com*.

Greenberg, A. (2014e). Silk Road 2.0 hacked using Bitcoin bug: All its funds stolen. *forbes. com*.

Greenberg, A. (2017). Global police spring a trap on thousands of dark web users. *wired. com*.

Grounded Theories Ltd (2016). Grounded theory online. *groundedtheoryonline. net*.

Gutman, L. M., Sameroff, A. J., and Cole, R. (2003). Academic growth curve trajectories from 1st grade to 12th grade: Effects of multiple social risk factors and preschool child factors. *Developmental psychology*, 39(4):777.

Hardy, R. A. and Norgaard, J. R. (2016). Reputation in the internet black market: an empirical and theoretical analysis of the deep web. *Journal of Institutional Economics*, 12(03):515–539.

Harfenist, E. and Turgeman, M. (2016). Dutch police open dark net site to spook vendors and buyers. *vocativ. com*.

Herbert, S. (1996). The normative ordering of police territoriality: making and marking space with the Los Angeles police department. *Annals of the Association of American Geographers*, 86(3):567–582.

Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual review of ecology and systematics*, 4(1):1–23.

Hollnagel, E. and Goteman, O. (2004). The functional resonance accident model. *Proceedings of cognitive system engineering in process plant*, 2004:155–161.

Hollnagel, E., Woods, D. D., and Leveson, N. (2007). *Resilience engineering: Concepts and precepts*. Ashgate Publishing, Ltd.

Horton-Eddison, M. and Di Cristofaro, M. (2017). Hard interventions and innovation in crypto-drug markets: The escrow example. *Policy Brief*.

Interpol (2018). Operations. *interpol. int*.

Iofciu, T., Fankhauser, P., Abel, F., and Bischoff, K. (2011). Identifying users across social tagging systems. In *ICWSM*.

Jeffries, A. (2013). Fbi seizes underground drug market Silk Road, owner indicted in new york. *theverge. com*.

Krebs (2015). Dark web's 'Evolution market' vanishes. *krebsonsecurity. com*.

Krebs (2017). Exclusive: Dutch cops on Alphabay 'refugees'. *krebsonsecurity. com*.

Kshetri, N. (2006). The simple economics of cybercrimes. *IEEE Security & Privacy*, 4(1):33–39.

Lacson, W. and Jones, B. (2016). The 21st century darknet market: Lessons from the fall of Silk Road. *International Journal of Cyber Criminology*, 10(1):40.

Lavorgna, A. (2016). How the use of the internet is affecting drug trafficking practices.

Leafly (2017). LSD strain. *leafly. com*.

Lee, J. P. and Antin, T. M. (2011). How do researchers categorize drugs, and how do drug users categorize them? *Contemporary drug problems*, 38(3):387–427.

Limited, F. M. (2016). GBP/USD – British Pound US Dollar. *uk. investing. com*.

Lorenzo-Dus, N. and Di Cristofaro, M. (2018). 'I know this whole market is based on the trust you put in me and i don't take that lightly': Trust, community and discourse in crypto-drug markets. *Discourse & Communication*, page 1750481318771429.

Ludwig, D., Walker, B., and Holling, C. S. (1997). Sustainability, stability, and resilience. *Conservation ecology*, 1(1).

Lund Research Ltd (2012). Self-selection sampling. *dissertation. laerd. com*.

Luo, Q. (2017). An exploratory investigation into the darknet marketplace discussion forum Agora. Master's thesis, UCL.

Luthar, S. S. (2006). Resilience in development: A synthesis of research across five decades.

Mac, R. (2013a). Lessons from Silk Road: Competing online drug site shuts down after security breach. *forbes. com*.

Mac, R. (2013b). Meet the Silk Road employee that the Dread Pirate Roberts alledgedly tried to murder. *forbes. com*.

Madni, A. M. and Jackson, S. (2009). Towards a conceptual framework for resilience engineering. *IEEE Systems Journal*, 3(2):181–191.

Magis, K. (2010). Community resilience: An indicator of social sustainability. *Society and Natural Resources*, 23(5):401–416.

Manchin, J. (2011). Manchin urges federal law enforcement to shut down online black market for illegal drugs. *manchin.senate.gov*.

Market, A. (2013). Atlantis Market. *facebook. com*.

Markopoulos, P., Xefteris, D., and Dellarocas, C. (2015). Manipulating reviews in dark net markets to reduce crime.

Martin, J. (2014). Lost on the Silk Road: Online drug distribution and the 'cryptomarket'. *Criminology & Criminal Justice*, 14(3):351–367.

Masten, A. S. (2001). Ordinary magic: Resilience processes in development. *American psychologist*, 56(3):227.

Masten, A. S., Best, K. M., and Garmezy, N. (1990). Resilience and development: Contributions from the study of children who overcome adversity. *Development and psychopathology*, 2(4):425–444.

Masten, A. S., Morison, P., Pellegrini, D., and Tellegen, A. (1992). 11 competence under stress: risk and protective factors. *Risk and protective factors in the development of psychopathology*, page 236.

McDonald, C. G. and Slawson, V. C. (2002). Reputation in an internet auction market. *Economic Inquiry*, 40(4):633–650.

Melnik, M. I. and Alm, J. (2002). Does a seller's ecommerce reputation matter? evidence from eBay auctions. *The journal of industrial economics*, 50(3):337–349.

Moore, T. and Clayton, R. (2007). Examining the impact of website take-down on phishing. In *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, pages 1–13. ACM.

Morselli, C., Décary-Hétu, D., Paquet-Clouston, M., and Aldridge, J. (2017). Conflict management in illicit drug cryptomarkets. *International Criminal Justice Review*, page 1057567717709498.

Motoyama, M., McCoy, D., Levchenko, K., Savage, S., and Voelker, G. M. (2011). An analysis of underground forums. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pages 71–80. ACM.

Mullin, J. (2014). Silk Road 2.0 infiltrated from the start sold 8m per month in drugs. *arstechnia. com*.

Munksgaard, R. and Demant, J. (2016). Mixing politics and crime–the prevalence and decline of political discourse on the cryptomarket. *International Journal of Drug Policy*, 35:77–83.

Munksgaard, R., Demant, J., and Branwen, G. (2016). A replication and methodological critique of the study "evaluating drug trafficking on the Tor network". *International Journal of Drug Policy*, 35:92–96.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Norris, F. H., Stevens, S. P., Pfefferbaum, B., Wyche, K. F., and Pfefferbaum, R. L. (2008). Community resilience as a metaphor, theory, set of capacities, and strategy for disaster readiness. *American journal of community psychology*, 41(1-2):127–150.

Nurmi, J., Kaskela, T., Perälä, J., and Oksanen, A. (2017). Seller's reputation and capacity on the illicit drug markets: 11-month study on the finnish version of the Silk Road. *Drug and alcohol dependence*, 178:201–207.

O'Neill, P. H. (2013). Cops may have just busted a major illegal gun dealer from the deep web. *dailydot. com*.

Ormsby, E. (2016). Silk Road: insights from interviews with users and vendors. *The Internet and drug markets*, pages 61–68.

Pastrana, S., Hutchings, A., Caines, A., and Buttery, P. (2018). Characterizing Eve: Analysing cybercrime actors in a large underground forum.

Paternoster, R. (1989). Absolute and restrictive deterrence in a panel of youth: Explaining the onset, persistence/desistance, and frequency of delinquent offending. *Social Problems*, 36(3):289–309.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Phelps, A. and Watt, A. (2014). I shop online–recreationally! internet anonymity and Silk Road enabling drug use in Australia. *Digital Investigation*, 11(4):261–272.

Pollack, H. A. and Reuter, P. (2014). Does tougher enforcement make drugs more expensive? *Addiction*, 109(12):1959–1966.

Przepiorka, W., Norbutas, L., and Corten, R. (2017). Order without law: Reputation promotes cooperation in a Cryptomarket for illegal drugs. *European Sociological Review*.

Redakcja (2015). [aktualizacja] wyciek bazy torepublic. *niebezpiecznik.pl*.

Reddit (2018). /r/ebay. *reddit.com*.

Reddit User (2013a). Should we be worried? *reddit.com*.

Reddit User (2013b). Warning: the Silk Road revealed its public IP. *reddit.com*.

Reddit User (2014a). [alpaca] we are re-opened! *reddit.com*.

Reddit User (2014b). Black Goblin obituary. *reddit.com*.

Reddit User (2014c). Darknet markets search engine. *reddit.com*.

Reddit User (2014d). ECC (eastcoastcollective) 'runs' cannabis road, both old and new versions. *reddit.com*.

Reddit User (2014e). I made it all pretty and stuff. *reddit.com*.

Reddit User (2014f). New market.. Exxtacy. *reddit.com*.

Reddit User (2014g). Onionshop hacked, all Bitcoins stolen (and i'm still in the db...). *reddit.com*.

Reddit User (2014h). [the Armory] new currency accepted: Darkcoin. *reddit. com*.

Reddit User (2015a). Agora to pause operations. *reddit. com*.

Reddit User (2015b). Amazon Dark - new market announcement. *reddit. com*.

Reddit User (2015c). The details of what happened at Kiss. *reddit. com*.

Reddit User (2015d). [PSA/article] Evo down, journalists incoming, watch what you post. *reddit. com*.

Reddit User (2016). Seized parcel help please. *reddit. com*.

Reddit User (2017). Darknetmarkets. *reddit. com*.

Reggiani, A., De Graaff, T., and Nijkamp, P. (2002). Resilience: an evolutionary approach to spatial economic systems. *Networks and Spatial Economics*, 2(2):211–229.

Rekšņa, T. et al. (2017). Complex network analysis of darknet black market forum structure. Master's thesis.

Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Experimental economics*, 9(2):79–101.

Richardson, G. E. (2002). The metatheory of resilience and resiliency. *Journal of clinical psychology*, 58(3):307–321.

Robins, G., Pattison, P., Kalish, Y., and Lusher, D. (2007). An introduction to exponential random graph (p*) models for social networks. *Social networks*, 29(2):173–191.

Schmidt, M. M., Sharma, A., Schifano, F., and Feinmann, C. (2011). "legal highs" on the net-evaluation of UK-based websites, products and product information. *Forensic science international*, 206(1):92–97.

Schnabel, Z. E. (1938). The estimation of total fish population of a lake. *The American Mathematical Monthly*, 45(6):348–352.

Scott, J. (2017). *Social network analysis*. Sage.

Seidman, E. and Pedersen, S. (2003). Holistic contextual perspectives on risk, protection, and competence among low-income urban adolescents. *Resilience and vulnerability: Adaptation in the context of childhood adversities*, pages 318–342.

Shapiro, S. S. and Wilk, M. B. (1965). An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611.

Silk Road Drugs (2014). US senator has called for a crackdown on the dark web. *silkroaddrugs. org*.

Soska, K. and Christin, N. (2015). Measuring the longitudinal Evolution of the online anonymous marketplace ecosystem. In *USENIX Security*, volume 15.

Soudijn, M. R. and Zegers, B. C. T. (2012). Cybercrime and virtual offender convergence settings. *Trends in organized crime*, 15(2-3):111–129.

Spitters, M., Klaver, F., Koot, G., and van Staalduinen, M. (2015). Authorship analysis on dark marketplace forums. In *Intelligence and Security Informatics Conference (EISIC), 2015 European*, pages 1–8. IEEE.

Stuck_In_the_Matrix (2015). I have every publicly available reddit comment for research. 1.7 billion comments 250 GB compressed. any interest in this? *reddit. com*.

Synopsis Inc (2014). The heartbleed bug. *heartbleed. com*.

Tarquin (2014). The Majestic Garden. *darkwebnews. com*.

Teddlie, C. and Yu, F. (2007). Mixed methods sampling: A typology with examples. *Journal of mixed methods research*, 1(1):77–100.

tfidf.com (2018). Term frequency – inverse document frequency. *tfidf. com*.

The Ross Ulbircht Legal Defense Effort (2018). Double standard. *freeross. org*.

The Ross Ulbricht Legal Defense Effort (2018). The Ross Ulbricht legal defense effort. *freeross. org*.

Tor Project (2017). Tor: Overview. *torproject. org*.

U.S. Attoney's Office (2015). Ross Ulbricht, AKA Dread Pirate Roberts, sentenced in Manhattan Federal Court to life in prison. *FBI. gov*.

U.S. Attorney's Office (2014a). Dozens of online 'dark markets' seized pursuant to forfeiture complaint filed in Manhattan Federal Court in conjunction with the arrest of the operator of Silk Road 2.0. *FBI. gov*.

U.S. Attorney's Office (2014b). Operator of Silk Road 2.0 website charged in manhattan central court. *FBI. gov*.

U.S. Attorney's Office (2014c). Operator of Silk Road 2.0 website charged in Manhattan Federal Court. *FBI. gov*.

U.S. Attorney's Office (2015). Ross Ulbricht, the creator and owner of the Silk Road website, found guilty in Manhattan Federal Court on all counts. *FBI. gov*.

U.S. District Court for the Southern District of New York (2013). United States of America v. Ross William Ulbricht. *documentcloud. org*.

U.S. District Court for the Western District of Washington at Seattle (2017). United States of America v. Brian Farrell. documentcloud.org.

Van Breda, A. D. et al. (2001). Resilience theory: A literature review. *Pretoria, South Africa: South African Military Health Service*.

Van Buskirk, J., Bruno, R., Dobbins, T., Breen, C., Burns, L., Naicker, S., and Roxburgh, A. (2017). The recovery of online drug markets following law enforcement and other disruptions. *Drug and alcohol dependence*, 173:159–162.

Van Buskirk, J., Naicker, S., Roxburgh, A., Bruno, R., and Burns, L. (2016a). Who sells what? country specific differences in substance availability on the Agora cryptomarket. *International Journal of Drug Policy*, 35:16–23.

Van Buskirk, J., Roxburgh, A., Bruno, R., Naicker, S., Lenton, S., Sutherland, R., Whittaker, E., Sindicich, N., Matthews, A., Butler, K., et al. (2016b). Characterising dark net marketplace purchasers in a sample of regular psychostimulant users. *International Journal of Drug Policy*, 35:32–37.

Van Buskirk, J., Roxburgh, A., Farrell, M., and Burns, L. (2014). The closure of the Silk Road: What has this meant for online drug trading? *Addiction*, 109(4):517–518.

Van Buskirk, J., Roxburgh, A., Naicker, S., and Burns, L. (2015). A response to Dolliver's "evaluating drug trafficking on the Tor network". *International Journal of Drug Policy*, 26(11):1126–1127.

Van Hout, M. C. and Bingham, T. (2013a). 'Silk Road', the virtual drug marketplace: A single case study of user experiences. *International Journal of Drug Policy*, 24(5):385–391.

Van Hout, M. C. and Bingham, T. (2013b). 'surfing the Silk Road': A study of users' experiences. *International Journal of Drug Policy*, 24(6):524–529.

Van Hout, M. C. and Bingham, T. (2014). Responsible vendors, intelligent consumers: Silk Road, the online revolution in drug trading. *International Journal of Drug Policy*, 25(2):183–189.

Van Wegberg, R., Tajalizadehkhoob, S., Soska, K., Akyazi, U., Ganan, C. H., Klievink, B., Christin, N., and van Eeten, M. (2018). Plug and prey? measuring the commoditization of cybercrime via online anonymous markets. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1009–1026. USENIX Association.

Wadsworth, E., Drummond, C., Kimergård, A., and Deluca, P. (2017). A market on both "sides" of the law: The use of the hidden web for the sale of new psychoactive substances. *Human Psychopharmacology: Clinical and Experimental*.

Walker, B., Holling, C. S., Carpenter, S., and Kinzig, A. (2004). Resilience, adaptability and transformability in social–ecological systems. *Ecology and society*, 9(2).

Walker, B., Kinzig, A., and Langridge, J. (1999). Plant attribute diversity, resilience, and ecosystem function: the nature and significance of dominant and minor species. *Ecosystems*, 2(2):95–113.

Wang, X., Peng, P., Wang, C., and Wang, G. (2018). You are your photographs: Detecting multiple identities of vendors in the darknet marketplaces.

Weisburd, D., Wyckoff, L. A., Ready, J., Eck, J. E., Hinkle, J. C., and Gajewski, F. (2006). Does crime just move around the corner? a controlled study of spatial displacement and diffusion of crime control benefits. *Criminology*, 44(3):549–592.

Wikipedia (2017). Tor mail. *en. wikipedia. org*.

Young, R., Zhang, L., and Prybutok, V. R. (2007). Hacking into the minds of hackers. *Information Systems Management*, 24(4):281–287.

Zetter, K. (2013). How the feds took down the Silk Road drug wonderland. *wired. com*.