



## City Research Online

### City, University of London Institutional Repository

---

**Citation:** Bracci, A. (2023). Blockchain based ecosystems: a complex systems approach. (Unpublished Doctoral thesis, City, University of London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

---

**Permanent repository link:** <https://openaccess.city.ac.uk/id/eprint/29882/>

**Link to published version:**

**Copyright:** City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

**Reuse:** Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

CITY, UNIVERSITY OF LONDON

---

# Blockchain based ecosystems: a complex systems approach

---

*Author:*

Alberto Bracci

*First Supervisor:* Dr. Andrea Baronchelli

*Second Supervisor:* Dr. Olalla Castro

*Examiners:*

Prof. Mark Broom

and

Dr. Federico Botta

*A doctoral dissertation submitted in fulfilment of the requirements for the  
degree of Doctor of Philosophy  
in the*

Department of Mathematics  
School of Science and Technology

February 17, 2023

# Contents

<b>List of Figures</b>	<b>5</b>
<b>List of Tables</b>	<b>20</b>
<b>1 Introduction</b>	<b>26</b>
<b>2 Background</b>	<b>32</b>
2.1 Where it all began: Bitcoin . . . . .	32
2.2 Dark Web Marketplaces . . . . .	34
2.3 Ethereum . . . . .	35
2.4 Data sources . . . . .	37
2.4.1 DWM listings . . . . .	37
2.4.2 Bitcoin transactions . . . . .	39
2.4.3 NFT transactions . . . . .	41
<b>3 Dark Web Marketplaces: reacting to the COVID-19 pandemic</b>	<b>43</b>
3.1 Data and methods . . . . .	44
3.1.1 Dark web marketplaces . . . . .	44
3.1.2 Twitter . . . . .	46
3.1.3 Wikipedia . . . . .	47
3.2 Results . . . . .	47
3.2.1 Categories of listings . . . . .	47
3.2.2 Time evolution of DWM listings and public attention . . . . .	51
3.2.3 Impact of COVID-19 on other listings . . . . .	54
3.3 Conclusions . . . . .	54
<b>4 Dark Web Marketplaces: adapting to the COVID-19 pandemic</b>	<b>57</b>
4.1 Data and methods . . . . .	58
4.2 Results . . . . .	60
4.2.1 Vaccine listings . . . . .	60
4.2.2 Other COVID-19 related products . . . . .	64
4.2.3 Listings with COVID-19 mentions . . . . .	66

4.3	Discussions and conclusion . . . . .	67
<b>5</b>	<b>Macroscopic properties of buyer-seller networks in online marketplaces</b>	<b>70</b>
5.1	Data . . . . .	72
5.1.1	Dark Web Marketplaces . . . . .	72
5.1.2	E-Commerce Platform . . . . .	72
5.2	Results . . . . .	72
5.2.1	Empirical properties of buyer-seller networks . . . . .	72
5.2.2	Modeling buyer-seller networks . . . . .	76
5.3	Discussion . . . . .	80
<b>6</b>	<b>The decentralized evolution of decentralization across fields: from Governance to Blockchain</b>	<b>82</b>
6.1	The pipeline . . . . .	83
6.1.1	Step 1: Data Collection . . . . .	84
6.1.2	Step 2: Hierarchical clustering . . . . .	84
6.1.3	Step 3: Knowledge flows . . . . .	85
6.2	Results . . . . .	86
6.2.1	The decentralized evolution of (de)centralization . . . . .	86
6.2.2	The case of Governance and Blockchain . . . . .	91
6.3	Discussion . . . . .	95
6.4	Code and data availability . . . . .	96
<b>7</b>	<b>Emergence and structure of decentralised trade networks around dark web marketplaces</b>	<b>97</b>
7.1	Data . . . . .	99
7.2	Results . . . . .	100
7.2.1	Large number of U2U transactions . . . . .	100
7.2.2	Behaviour of the U2U network . . . . .	101
7.2.3	U2U network evolution . . . . .	103
7.3	Discussion and Conclusion . . . . .	107
<b>8</b>	<b>Heterogeneous rarity patterns drive price dynamics in NFT collections</b>	<b>109</b>
8.1	Background, data and methods . . . . .	111
8.1.1	Glossary of key terms . . . . .	111
8.1.2	Dataset . . . . .	111
8.1.3	Rarity . . . . .	112
8.2	Results . . . . .	113
8.2.1	Market Growth . . . . .	113
8.2.2	Quantifying rarity . . . . .	115
8.2.3	Rarity and market performance . . . . .	116
8.2.4	Rarity and return on investment . . . . .	118
8.3	Discussion . . . . .	119

<b>9</b>	<b>Conclusions</b>	<b>122</b>
<b>A</b>	<b>Appendix to chapter 3</b>	<b>127</b>
A.1	Data pre-processing . . . . .	127
A.2	Examples of listings related to COVID-19 in dark web marketplaces . . . . .	129
A.2.1	COVID-19 specific listings . . . . .	129
A.2.2	COVID-19 mentions . . . . .	130
A.3	Timeline of the COVID-19 pandemic . . . . .	131
A.4	Additional material . . . . .	132
<b>B</b>	<b>Appendix to chapter 4</b>	<b>137</b>
B.1	Examples of detected listings . . . . .	137
B.2	DWMs offering COVID-19 vaccines . . . . .	143
B.3	Additional material . . . . .	144
<b>C</b>	<b>Appendix to chapter 5</b>	<b>149</b>
C.1	Computing the Memory Kernel . . . . .	149
C.2	Sampling of Product Markets . . . . .	150
C.3	Model parameters estimation . . . . .	150
C.4	Model simulation . . . . .	151
C.5	Additional Figures . . . . .	151
C.6	Additional Tables . . . . .	163
<b>D</b>	<b>Appendix to chapter 6</b>	<b>166</b>
D.1	Dataset . . . . .	166
D.2	Hierarchical clustering . . . . .	169
D.2.1	Keywords annotation . . . . .	169
D.2.2	Results of the hSBM . . . . .	173
D.3	The importance of governance and blockchain in the history of (de)centralization	175
D.3.1	Influence of governance from and on other clusters . . . . .	177
<b>E</b>	<b>Appendix to chapter 7</b>	<b>178</b>
E.1	Additional methods . . . . .	178
E.2	General statistics of the 40 DWMs under consideration . . . . .	179
E.3	Detection of stable pairs in temporal and directed networks . . . . .	183
E.3.1	Temporal and undirected networks . . . . .	183
E.3.2	Temporal and directed networks . . . . .	184
E.4	Additional simulations . . . . .	186
<b>F</b>	<b>Appendix to chapter 8</b>	<b>191</b>
F.1	OpenSea market mechanisms . . . . .	191
F.2	Generative art mechanisms . . . . .	192
F.3	List of collections . . . . .	192
F.4	Rarity score distributions . . . . .	193

F.5	Additional figures . . . . .	193
F.5.1	Rarity rank results . . . . .	196
F.5.2	Currency robustness check . . . . .	198
F.5.3	Time robustness check . . . . .	199
F.5.4	Tails robustness check . . . . .	202
<b>Bibliography</b>		<b>204</b>

# List of Figures

2.1	<b>Example of a DWM listing.</b> Screenshot of a chloroquine listing in the DarkBay/Dbay marketplace, where we highlight some of its salient attributes. Among the attributes considered in this work and shown in Table A.1, “Time” and “Marketplace name” attributes are not present in this screenshot, while the “Quantity” attribute is not fixed by the vendor. . . . .	38
2.2	<b>Identification of real entities in the Blockchain.</b> End goal of Bitcoin transactions clustering techniques: mapping a series of Bitcoin addresses to real entities. In this example, an address sends Bitcoins to another address. Thanks to the identification process, the two addresses are associated with two real entities. The Bitcoin transaction between the two entities becomes traceable and transparent. . . . .	40
3.1	<b>COVID-19 related listings: price and shipping geography.</b> (a) Box plot of the distribution of listing prices for each COVID-19 category. The box ranges from the lower to the upper quartile, with the horizontal line indicating the median. The whiskers extend up to the 5 <sup>th</sup> and 95 <sup>th</sup> percentiles respectively. The dots represent outliers. (b) Shipping information in COVID-19 specific listings. Note that 545 (or 71.1%) of these listings did not provide any shipping information. . . . .	49
3.2	<b>COVID-19 related listings: DWMs and vendors.</b> (a) Word cloud for “Listing title” in COVID-19 specific listings. (b) Category breakdown of COVID-19 specific listings in the DWM that offered them. (c) Fraction of vendors selling at least one COVID-19 specific listing. (d) Vendor specialisation. Most vendors responsible for at least one COVID-19 specific listing also sell other listings, and in greater number. . . . .	50

3.3	<b>Longitudinal analysis of DWM activity.</b> (a) Seven-days rolling average of active listings mentioning COVID-19 and COVID-19 specific listings. (b) Seven-days rolling average of the observed COVID-19 specific listings in the <i>medicines</i> and <i>PPE</i> categories. Black dashed vertical lines in panels (a) and (b) corresponded to significant COVID-19 world events, see Appendix A.3. (c) Seven-days median price with 95% confidence interval for COVID-19 specific listings. (d) Seven-days median price with 95% confidence interval for active COVID-19 specific listings in the <i>PPE</i> and <i>medicines</i> categories. . . . .	52
3.4	<b>DWMs and public attention.</b> (a)-(c) Seven-days rolling average of active listings selling <i>PPE</i> , together with the time evolution of the number of tweets referring to masks and of visits in the relative Wikipedia page visits. (b)-(d) Similar comparison as in panels (a)-(c) but considering active listings of hydroxychloroquine, chloroquine, and azithromycin. Black dashed vertical lines in panels (a) and (b) mark significant events related with COVID-19, see Appendix A.3. See Appendix A.4 for panels (a) and (b) with a linear y-axis. . . . .	53
3.5	<b>The indirect impact of the COVID-19 pandemic.</b> Percentage of all active listings mentioning the themes lockdown, delay and sales in panels (a), (b), (c), respectively. (d) Percentage of active listings in panels (a), (b), (c) that mentioned also COVID-19 in their listings. Black dashed vertical lines in panels (a), (b), and (c) corresponded to major COVID-19 events, see Appendix A.3. . . . .	55
4.1	<b>DWMs and COVID-19 vaccines.</b> (a) Number of unique listings offered in each DWM. “BB. house” stands for Big brother house, while “Canadian HQ.” to The Canadian Headquarters. (b) Violin plot of the prices in USD at which vaccines were offered, showing the distribution of prices for the three categories. (c) Number of listings indicating where vaccines are declared to be shipped from and to. “N.A.” stands for not applicable and “Russia” for Russia and Eastern neighbouring countries. (d) Number of vendors offering a vaccine in a given category. Only the lower triangle of the matrix is shown because it is symmetric, where its diagonal represents vendors offering only listings in that category. . . . .	61
4.2	<b>Temporal evolution of COVID-19 vaccine listings.</b> (a) Cumulative number of listings over time in the three categories considered. (b) Average price over time in the same three categories, computed with a 90-days moving window. Vertical dashed black lines represent relevant pandemic events. . . . .	63

4.3	<b>COVID-19 related products.</b> (a) Break-down of COVID-19 related products by category and market. (b) Boxplot representing the price distribution of listings in each category. Horizontal lines represent the median value, box ends the first and third quartiles, and whiskers minimum and maximum values, respectively. (c) Number of listings indicating where COVID-19 related products are declared to be shipped from and to. “N.A.” stands for not applicable and “Russia” for Russia and Eastern neighbouring countries. (d) Number of vendors offering a COVID-19 related product in a given category. Only the lower triangle of the matrix is shown because it is symmetric, where its diagonal represents vendors offering only listings in that category. . . . .	65
4.4	<b>Characterization of COVID-19 mentions.</b> (a) UMAP representation of doc2vec embeddings. hDBSCAN clustering finds 13 meaningful categories covering COVID-19 related products and all major drugs sold on DWMs. (b) 30-day rolling average of active listings in 4 categories of listings mentioning COVID-19: <i>thc</i> , <i>psychedelics</i> , <i>cocaine</i> , and <i>mushrooms</i> . (c) 30-day rolling average of fraction of previously identified drugs listings mentioning 3 different COVID-19 related themes: lockdowns, shipping delays and sales, with key pandemic events highlighted with vertical dashed black lines. . . . .	68
5.1	<b>Online marketplaces show strikingly similar patterns according to different aggregated quantities.</b> Top (a to d): Distributions of 4 analyzed users aggregate quantities: money and number of transactions both sent and received. Bottom (e to h): Distributions of 4 analyzed users temporal quantities: the inter-event time (time between successive transactions) and the lifetime (time between first and last transaction) both measured in days. Each blue line represents one DWM, the black line is the distribution built aggregating all DWMs together, each dashed red line represents one e-commerce platform market, while the dashed yellow line is the distribution aggregating all e-commerce markets. Similar patterns are observed between different markets in the same platform, but also across regulated and unregulated online marketplaces. . . . .	73
5.2	<b>Buyer memory affects their purchase decisions.</b> a) Normalised buyer entropy distribution for each e-commerce market (red), and the whole e-commerce platform (yellow), excluding users with zero entropy (mostly users with one transaction, see Fig. C.1) for visual clarity. The fat-tailed distributions span the full $[0, 1]$ range, with most buyers almost equally buying from multiple sellers. b) $P(n)$ is the probability to buy from a new seller after a buyer has already bought from $n$ different ones. Each degree class $k_{min} \leq k \leq 2 \cdot k_{min} - 1$ is rescaled according to the fitted value of $c$ and $\beta$ (see Tab. C.1 for the values), with Eq.2 (dashed line) well reproducing the memory effect on the buyers’ behaviour: the more sellers they try, the less likely they are to buy from a new one. . . . .	75

5.3	<b>Empirical activity distributions.</b> a): Activity distribution for all e-commerce markets in red, activity distribution of all aggregated markets in yellow. b): Activity distribution for all DWMs in blue, activity distribution of all aggregated DWMs in black. This plot is purely descriptive as the DWM activity distributions are not used in the rest of the analysis. . . . .	77
5.4	<b>The model reproduces different properties of buyer-seller networks.</b> Each row corresponds to a different market (see Fig. C.3-C.6 for other markets), whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships. . . . .	78
5.5	<b>Model reproduces the temporal evolution of the top sellers degree distribution.</b> Top (a to c): temporal evolution of the degree distribution of the top 50 (a), 100 (b) and 200(c) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5 <sup>th</sup> to 97.5 <sup>th</sup> percentiles. Results are shown for one product market, all other markets are shown in Fig. C.7-C.10. Bottom (d to f): Distribution of the distance between the empirical and model(s) median degree of the top 50(left), 100(center) and 200(right) sellers, for all product markets and time steps, and the three considered models. Vertical lines represent the distributions median, showing that the model median is always smaller than the alternatives. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism. . . . .	79

- 6.1 **The rising interest of academic literature towards *(de)centralization*.** (a) Fraction of papers mentioning *(de)centralization* in the Semantic Scholar corpus. (b) Fraction of authors mentioning *(de)centralization* in the Semantic Scholar corpus. Both fractions have been steadily increasing since the 1950s, showing growing interest in the topic. (c) Number of papers in the clusters at the 3<sup>rd</sup> level of the hierarchy in each year. Clusters are ordered respecting the hierarchical network on the left, in which node and link sizes are proportional to the total number of papers in the related cluster. In the heatmap, white lines individuate clusters belonging to the same cluster at the 2<sup>nd</sup> level of the hierarchy, while red lines divide different clusters at the 1<sup>st</sup> level. The representative keyword of the clusters at level 3 is reported in bold in the respective rows, while all the specific keywords identified at the 4<sup>th</sup> hierarchical level are shown within brackets. Clusters with less than 500 papers in total are not shown in the figure. . . . . 88
- 6.2 **Temporal evolution of the influence between clusters.** Average knowledge flows  $K_{a \rightarrow b}(T)$  from each cluster  $a$  to each cluster  $b$  at level 3 in the period  $T = 1970-1989$  (a),  $T = 1990-2007$  (b) and  $T = 2008-2020$  (c), represented by the different colors according to the colorbar in the bottom left part of the figure. A representative keyword for each cluster is reported on the bottom right part of the figure. White lines denote clusters belonging to the same 2<sup>nd</sup> level cluster, whereas red lines mark different branches at the 1<sup>st</sup> level. In the first period, little to no communication is happening between different clusters. In recent years, more communication happens inside the same 2<sup>nd</sup> level cluster, and towards the middle branch (cluster 11). However, little communication happens between the two other branches, roughly representing the STEM and social sciences communities respectively. . . . . 90
- 6.3 **Ranking the influence of Blockchain and Governance in the *(de)centralization literature*.** (a) Ranking in time of the influence coming from a cluster in the 3<sup>rd</sup> hierarchical level, computed on the average knowledge flow  $K_{a \rightarrow \bullet}(Y)$  from papers in cluster  $a$  published in the year  $Y$  towards all other future papers. (b) Ranking in time of the influence to a cluster in the 3<sup>rd</sup> hierarchical level, computed on the average knowledge flow  $K_{\bullet \rightarrow a}(Y)$  from papers published in the year  $Y$  towards future papers in cluster  $a$ . The blockchain cluster, highlighted in blue, has become a central actor in the recent literature on *(de)centralization*, supplanting the governance cluster, highlighted in red. . . . 92

6.4	<b>Influences between Blockchain and the other clusters on <i>(de)centralization</i>.</b>	
	(a) Change in the ranking of the most influenced clusters by Blockchain between its early period (2008-2014), its middle period (2015-2018), and its late period (2019-2020), calculated using the average knowledge flows $K_{a \rightarrow b}(T)$ , where $T$ is the selected period, and $a$ is fixed to be Blockchain. (b) Change in the ranking of the clusters having most influenced the Blockchain literature (same periods as in the previous panel). In both cases, if $K_{a \rightarrow b}(T) = 0$ , we print a circle in the corresponding gray node and use a lighter color in the respective link. Moreover, we print a star when $0.01 < K_{a \rightarrow b}(T) \leq 0.1$ , and two stars when $K_{a \rightarrow b}(T) > 0.1$ . . . . .	94
7.1	<b>Ego and full networks.</b> (a) Schematic representation of an ego network surrounding a dark web marketplace (“DWM”, in red). The DWM interacts with its users (“U”, in black), which make user-to-user (U2U) pairs, represented with arrows and their respective users. (b) Multiple ego networks may be aggregated to form the full network. . . . .	98
7.2	<b>User-DWM and U2U transactions.</b> (a) Total number of users interacting with a DWM against the total number of them forming U2U transactions. The dotted line corresponds to the result of a fitted power law function. (b) Trading volume in dollars sent to a DWM compared with the total trading volume in its surrounding U2U transactions. The dashed line is the bisector and allows to easily compare the two trading volumes. (c) Total monthly trading volume sent to all DWMs and exchanged in all unique U2U pairs. We do not include the trading volume received from DWMs because it is equivalent to the volume sent to DWMs. . . . .	100
7.3	<b>U2U network.</b> The U2U network is formed by the entire set of interacting users (black and gray arrows with their respective users). Using the evolving activity-driven model [1], U2U pairs are divided in either stable (black arrows and respective users) or unstable (gray arrows and respective users). . . . .	102
7.4	<b>Role of users forming stable U2U pairs.</b> (Main) PDFs of trading volume that users exchange with any DWMs. (Inset) PDFs of time spent by users on any DWMs. These distributions are explored for each of the 40 DWMs under consideration in Figure E.4 and E.5, respectively. Vertical lines represent median values of the respective distributions. . . . .	102
7.5	<b>Resilience of stable U2U pairs after DWMs closure.</b> Trading volume of U2U pairs surrounding active DWMs. (Main) U2U pairs meet who met inside aa DWM. (Inset) U2U pairs meet outside them. Curves indicate the median value while bands represent the 95% confidence interval. Day zero corresponds to the day when the market closed. Negative and positive numbers indicate the days prior and after the closure, respectively. Only the 33 DWMs closed are considered in the analysis. . . . .	105
7.6	<b>Temporal evolution of stable pairs.</b> (a) Monthly number of new stable U2U pairs created. (b) Monthly trading volume of stable U2U pairs. . . . .	106

8.1	<b>Illustration of rarity in a collection.</b> Demonstration of the visual difference between rare and common NFTs using the example of CryptoPunks. CryptoPunk #2547 (on the right) is the least rare, as it has traits that appear frequently in the collection (i.e., the bandanna and the earring). CryptoPunk #8348 is the rarest in the collection, mostly since it is the only one with seven non-null attributes. Rarity scores are not normalised. After normalisation, the total rarity score for punk #2547 is zero, while the one for #8348 is 100 (min and max of the collection, respectively). In the bottom left corner, we show, as an example, the rarity score of traits associated to the “Facial Hair” attribute. . . . .	110
8.2	<b>Characteristics of collections.</b> The probability distribution of a) the total number of NFTs within the collection; b) the total number of attributes in the collection; c) the total number of traits in the collection. . . . .	112
8.3	<b>The collectible market boom.</b> 30-day rolling average of the daily number of primary sales (blue line), secondary sales (orange line) and new buyers (i.e., new wallet addresses buying an NFT for the first time, green line). Inset: 30-day rolling average of the daily trade volume (in USD). . . . .	113
8.4	<b>Cumulative distribution of key market properties across collections.</b> a) Total number of sales per collection. b) Total trade volume per collection. c) Median sale price per collection. In inset: Distribution of each market property. . . . .	114
8.5	<b>Rare NFTs are scarce within a collection.</b> a-c) Distribution of the rarity score of the NFTs within three collections: CryptoPunks (a), Bored Ape Yacht Club (b), and World of Women (c). d) Violin plot of the Spearman Rank correlation computed between the rarity score and the number of NFTs with that score. 96% of the collections have a Spearman rank $r \leq 0$ (black dashed lines). . . . .	116
8.6	<b>Rare NFTs have a higher financial value and circulate less on the marketplace.</b> Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity $(100-q)$ , where $q$ is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. . . . .	118
8.7	<b>High rarity leads to higher returns, and a lower chance of a negative return.</b> a) Median return in USD by rarity quantile. Inset: median return against the quantity $(100-q)$ , where $q$ is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. . . . .	119

A.1	<b>Probability distribution function (Pdf) for the number of listings per vendor.</b> The power law fit results in an exponent of -2.3. In inset, the histogram of the number of listings per vendor, with a vertical line showing the 80 <sup>th</sup> percentile. . . . .	133
A.2	<b>Time evolution of the active COVID-19 specific listings in the <i>guides on scamming</i> category.</b> (a) Seven-days rolling average of these observed listings at a given time. Black dashed vertical lines corresponded to significant COVID-19 world events, see Appendix A.3. (b) Seven-days median price with 95% confidence interval for these observed listings. . . . .	134
A.3	<b>Wikipedia page visits for pages relative to (a) <i>PPE</i>, (b) hydroxychloroquine, chloroquine and azithromycin.</b> Number of tweets mentioning (c) <i>PPE</i> , (d) hydroxychloroquine, chloroquine and azithromycin. Panels (a) and (b) corresponds to Figure 3.4(a) in the main text, while panels (c) and (d) to Figure 3.4(b). The main difference between these panels and Figure 3.4(a) and (b) is the linear scale on y axis. . . . .	134
B.1	<b>Pfizer/BioNTech vaccine offered on Invictus.</b> Screenshots of a listing in the <i>approved vaccines</i> category offering the Pfizer/BioNTech vaccine at \$500 on the Invictus marketplace. We removed the contact information of the vendor, who invites the potential customer to have a direct contact. The screenshot was taken on February 6, 2021. . . . .	137
B.2	<b>Proof of vaccination offered on Hydra.</b> Proof of vaccination offered at 4,000 rubles (55 USD) detected on the Hydra marketplace. The original language of this listing is Russian and its translation in the English language is given in Table B.2. The screenshots were taken on February 6, 2021. . . .	138
B.3	<b>Summary of key events related to DWMs and COVID-19.</b> Availability of listings offering vaccines on dark web marketplaces (top), together with main COVID-19 related events of the vaccination campaign (bottom). . . . .	144
B.4	<b>Lifetime of listings on a DWM.</b> Number of days during which listings were active on a DWM. . . . .	145
B.5	<b>Categories of vaccines offered on DWMs.</b> Number of DWMs offering a vaccine in a given category. Only the lower triangle of the matrix is shown because it is symmetric, where its diagonal represents vendors offering only listings in that category. . . . .	145
B.6	<b>Vendor statistics.</b> Histograms representing the number of vendors offering a certain amount of vaccines listings, in panel (a), and the number of vendors trading in a given amount of DWMs, in panel (b). . . . .	145
B.7	<b>Price of COVID-19 <i>approved vaccines</i>.</b> Boxplots of the prices in USD at which vaccines were offered. (a) Price of listings in the three categories considered. (b) Focus on the listings offering approved vaccines. “J&J” stands for Johnson&Johnson. Horizontal lines represent the median value, box ends the first and third quartiles, and whiskers minimum and maximum values, respectively. . . . .	146

B.8	<b>Temporal evolution COVID-19 approved vaccines.</b> (a) Cumulative number of listings over time. (b) Average price over time, computed with a 90-days moving window. “J&J” stands for Johnson&Johnson. . . . .	146
B.9	<b>Time evolution of products mentioning COVID-19.</b> Number of active listings in time in each product category, according to the clustering described in the main text. . . . .	147
B.10	<b>Time evolution of fraction of all listings mentioning COVID-19 related themes.</b> . . . . .	148
C.1	<b>Most buyers with zero entropy only have done one transaction.</b> Histogram of percentage of buyers doing just one transaction for each e-commerce product market, among those with zero entropy. In most markets the percentage is well above 90%. . . . .	152
C.2	<b>The entropy distribution is significant against a null model.</b> (a) Buyer entropy distribution for each market (red) and all markets (yellow) in the e-commerce datasets. (b) Same distributions build reshuffling the transaction links, such that buyers have the same number of transactions, but with random sellers. The distributions with the reshuffled data show considerably higher entropy, meaning that the buyers in the dataset show non-random effects of memory. . . . .	152
C.3	<b>Model simulations for different markets - final distributions - Markets 3 to 9.</b> Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships. . .	153
C.4	<b>Model simulations for different markets - final distributions - Markets 10 to 16.</b> Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships. . .	154

C.5	<b>Model simulations for different markets - final distributions - Markets 17 to 23</b> Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships. . .	155
C.6	<b>Model simulations for different markets - final distributions - Markets 24 to 28</b> Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships. . .	156
C.7	<b>Model simulations for different markets - temporal evolution - Markets 2 to 8</b> Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5 <sup>th</sup> to 97.5 <sup>th</sup> percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism. . .	157
C.8	<b>Model simulations for different markets - temporal evolution - Markets 9 to 15</b> Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5 <sup>th</sup> to 97.5 <sup>th</sup> percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism. . .	158
C.9	<b>Model simulations for different markets - temporal evolution - Markets 16 to 22</b> Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5 <sup>th</sup> to 97.5 <sup>th</sup> percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism. . .	159

C.10	<b>Model simulations for different markets - temporal evolution - Markets 23 to 28</b>	Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5 <sup>th</sup> to 97.5 <sup>th</sup> percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism. . .	160
C.11	<b>Dark Web Marketplaces</b>	Duration of each market in our dataset, color coded by the number of transactions involving each marketplace. Each market is live at least for 180 days and averages at least 20'000 transactions per day.	161
C.12	<b>E-Commerce Platform Markets</b>	a): Plot of the number of transaction per each market of the e-commerce platform data. b): number of buyers and sellers in each market of the e-commerce platform data. . . . .	162
D.1	Comparison of the increase in number of papers <b>(a)</b> and authors <b>(b)</b> between the whole S2AG corpus and the <i>(de)centralization</i> dataset. Exponential fits are shown as dashed line, with the fitted values between brackets in the legend.		167
D.2	<b>Keywords annotation for Blockchain.</b> <b>(a)</b> Number of papers in time in each field of study for the Blockchain cluster. <b>(b)</b> Hierarchical topic tree, highlighting the top 20 words of the 10 most significant topics in the Blockchain cluster according to the normalized mixture proportion (in color). This information is used to aid the keywords annotation manual procedure. . . . .		170
D.3	<b>Keywords annotation for Governance.</b> <b>(a)</b> Number of papers in time in each field of study for the Governance cluster. <b>(b)</b> Hierarchical topic tree, highlighting the top 20 words of the 10 most significant topics in the Governance cluster according to the normalized mixture proportion (in color). This information is used to aid the keywords annotation manual procedure. . . .		171
D.4	<b>Results of hSBM at the 4<sup>th</sup> hierarchical level.</b> Heatmap of the number of papers in each cluster at level 4, together with the manually annotated keywords and the hierarchical tree resulting from the clustering algorithm. .		173
D.5	<b>Visualization of the results of the hSBM algorithm.</b> Bipartite network of documents and words, where nodes and links are colored based on the document cluster (at hierarchical level 3) they start from. Multiple instances of the same word are considered, one for each document the word is in. Word nodes within a topic at level 3 are ordered as the document clusters for visualization purposes. On top of the bipartite network, the hierarchy of topics and document clusters are represented on the word and document side respectively, shown as a network of square gray nodes. On the bottom of the figure, the legend with the 10 most frequent words of each topic is displayed.		174

D.6	<b>Blockchain and Governance are the most productive clusters.</b> (a) Number of papers in time for the Blockchain cluster, the Governance cluster and the full <i>(de)centralization</i> dataset. (b) Rank by number of papers in time for different clusters, highlighting Blockchain and Governance. Both plots show the central role of these two clusters in the literature on <i>(de)centralization</i> , and how they have exchanged roles, with Blockchain becoming the most productive field in recent years. . . . .	175
D.7	<b>Influence of Governance and Blockchain from and on other clusters in time.</b> Average knowledge flow towards other clusters in the future in time from a cluster (a) or to a cluster (b), the latter being either Blockchain (in blue) or Governance (in red), while in gray the average over the <i>(de)centralization</i> dataset is plotted. . . . .	176
D.8	<b>The influence from/on Governance with regards to the rest of the <i>(de)centralization</i> literature.</b> (a) Change of the ranking of the clusters most influenced by the Governance literature between its early period (1950-1980), its middle period (1981-1990), and its late period (1991-2000), calculated using the average knowledge flows $K_{a \rightarrow b}(T)$ , where $T$ is the selected period, and $a$ is fixed to Governance. (b) Change of the ranking of the most influential clusters on the Governance literature between its early period (1950-1980), its middle period (1981-1990), and its late period (1991-2000), calculated using the average knowledge flows $K_{a \rightarrow b}(T)$ , where $T$ is the selected period, and $b$ is fixed to Governance. In both cases, if $K_{a \rightarrow b}(T) = 0$ , we print a circle in the corresponding gray node and use a lighter color in the respective link. Moreover, we print a star when $0.01 < K_{a \rightarrow b}(T) \leq 0.1$ , and two stars when $K_{a \rightarrow b}(T) > 0.1$ . . . . .	177
E.1	<b>Lifetime of DWMs in our dataset.</b> Time interval between the first and last transaction of each DWM. A total of 17 DWMs participated in at least one transactions in either 2020 or 2021, while 23 closed before 2020. . . . .	181
E.2	<b>Key distributions of the full network.</b> Probability distribution function (pdf) about the number of transactions of each pair of entities (a)-(d)-(g), their trading volume (b)-(g)-(h), and their lifetime computed as time difference between their last and first transaction (c)-(f)-(i). . . . .	182
E.3	<b>Statistics of U2U pairs.</b> (a) Number of stable and non-stable U2U pairs with a given number of transactions. (b) Number of stable and non-stable U2U pairs with a given trading volume. Vertical lines represent median values of the respective distributions. . . . .	186
E.4	<b>Evaluation of the time users spent on a DWM.</b> It extends Figure 7.4(inset) in the main text by considering each individual DWM. Statistical tests are carried using the two-sided Kolmogorov-Smirnov test and results are available in Table E.3. Vertical lines represent median values of the respective distributions. . . . .	187

E.5	<b>Evaluation of the total trading volume users exchange with a DWM.</b> It extends Figure 7.4 in the main text by considering each individual DWM. Statistical tests are carried using the two-sided Kolmogorov-Smirnov test and results are available in Table E.3. Vertical lines represent median values of the respective distributions. . . . .	188
E.6	<b>DWMs where users meet.</b> Scatter plot of the number of pairs of users that meet inside each of the 40 DWMs considered versus the total volume sent to the DWM. . . . .	190
E.7	<b>Trading volume of stable U2U pairs during COVID-19.</b> 28-days moving average of trading volume between users who met inside a DWM (a) and outside any DWMs (b). (c) Yearly growth relative to the same day of 2019. Vertical lines represent the dates of Wuhan (Jan 23, 2020) and Italy (March 3, 2020) lockdowns. . . . .	190
E.8	<b>Formation of new stable U2U pairs during COVID-19.</b> 28-days moving average of new stable U2U pairs started between U2U pairs who met inside a DWM (a) and outside any DWMs (b). (c) Yearly growth relative to the same day of 2019. Vertical lines represent the dates of Wuhan and Italy lockdowns. . . . .	190
F.1	<b>Collectible NFTs minted over time.</b> Distribution of the collectible NFTs considered in this analysis minted over time. . . . .	194
F.2	<b>Distribution of the log-normal distribution characteristic parameter <math>\mu</math>.</b> Distribution of the log-normal distribution parameter $\mu$ (blue line), and its average value across collections (red dashed line). The log-normal distribution $\ln(X) \sim \mathcal{N}(\mu, \sigma^2)$ captures the distribution of rarity for 90% of collections. . . . .	194
F.3	<b>Fitting of the rarity distribution.</b> Distribution of the rarity score of the NFTs within several collections included in the dataset (blue dots), along with the best distribution fit computed using Maximum Likelihood Estimation and Akaike Information Criterion [2] (orange line). . . . .	195
F.4	<b>Trait Rarity Rank.</b> a-c) Distribution of the trait rarity rank of the NFTs within three collections: CryptoPunks (a), Bored Ape Yacht Club (b), and World of Women (c). d) Violin plot of the Spearman Rank correlation computed between the rarity rank and the number of NFTs with that rank. . . . .	196
F.5	<b>Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis with the rarity rank.</b> Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity $(100-q)$ , where $q$ is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. The NFTs are aggregated by quantile depending on their rarity rank, i.e the first quantile represents the rarest NFTs within the collection. . . . .	197

F.6	<b>High rarity leads to higher returns, and a lower chance of a negative return - Analysis with the rarity rank.</b> a) Median return in USD by rarity quantile. Inset: median return against the quantity $(100-q)$ , where $q$ is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. The NFTs are aggregated by quantile depending on their rarity rank, i.e the first quantile represents the rarest NFTs within the collection. . . . .	197
F.7	<b>Rare NFTs have a higher financial value and circulate less on the marketplace - Price in ETH.</b> Median sale price in ETH (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity $(100-q)$ , where $q$ is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. . . . .	198
F.8	<b>High rarity leads to higher returns, and a lower chance of a negative return - Price in ETH.</b> a) Median return in ETH by rarity quantile. Inset: median return against the quantity $(100-q)$ , where $q$ is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in ETH by rarity quantile. . .	198
F.9	<b>Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis on Q3 2021.</b> Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity $(100-q)$ , where $q$ is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. This analysis only takes into consideration the sales happening during Q3 2021. . .	199
F.10	<b>High rarity leads to higher returns, and a lower chance of a negative return - Analysis on Q3 2021.</b> a) Median return in USD by rarity quantile. Inset: median return against the quantity $(100-q)$ , where $q$ is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. This analysis only takes into consideration the sales happening during Q3 2021. . . . .	200

F.11	<b>Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis on Q4 2021.</b> Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity $(100-q)$ , where $q$ is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. This analysis only takes into consideration the sales happening during Q4 2021. . . . .	200
F.12	<b>High rarity leads to higher returns, and a lower chance of a negative return - Analysis on Q4 2021.</b> a) Median return in USD by rarity quantile. Inset: median return against the quantity $(100-q)$ , where $q$ is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. This analysis only takes into consideration the sales happening during Q4 2021. . . . .	201
F.13	<b>Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis without the rarest and least rare NFTs.</b> Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity $(100-q)$ , where $q$ is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. This analysis was performed after discarding the 10% rarest and least rare NFTs from each collection. . . . .	202
F.14	<b>High rarity leads to higher returns, and a lower chance of a negative return - Analysis without the rarest and least rare NFTs.</b> a) Median return in USD by rarity quantile. Inset: median return against the quantity $(100-q)$ , where $q$ is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. This analysis was performed after discarding the 10% rarest and least rare NFTs from each collection. . . . .	203

# List of Tables

3.1	<b>Categories used to classify the selected COVID-19 dataset.</b> The first five categories constitute COVID-19 specific listings, while the last one, called COVID-19 <i>mentions</i> , includes listings mentioning one of the keywords in Table A.2 without selling actual COVID-19 specific listings. . . . .	45
3.2	<b>Information on data related to each DWM.</b> Number of days each marketplace was crawled, the number of unique listings, all and COVID-19 specific, and the number of listing observations, all and COVID-19 specific. CanadianHQ indicates The Canadian HeadQuarters marketplace. . . . .	46
3.3	<b>Summary statistics for the considered categories of listings.</b> For each category, we included the number of unique listings, observations, and vendors. If the same vendor sold listings in different categories, we counted it as one unique vendor. . . . .	48
4.1	<b>Search of COVID-19 vaccines.</b> Keywords used to pre-select vaccine listings from the original dataset. Words are truncated to include different suffixes (e.g., vaccin yields vaccine, vaccination, vaccinate, etc.) . . . . .	59
4.2	<b>COVID-19 products related keywords.</b> Keywords used to pre-select listings selling COVID-19 related products before their manual annotation, organised by category. . . . .	59
7.1	<b>Formation mechanism of stable U2U pairs.</b> We compare the time in which the first transaction between a pair of users occur with the time in which these users interact with the same DWM. Each row in the figure indicates a possible temporal sequence, which we classify in two groups: users who met outside the DWM (first two columns) and users who met inside the DWM (last column). . . . .	104
7.2	<b>Nomenclature.</b> Definitions of all groups the users are divided to based on their behaviour. Number of users in each group is given in the last column. .	104
A.1	<b>Selected attributes of the listings under consideration,</b> along with a brief explanation of their respective purposes. . . . .	127
A.2	<b>Keywords used to sample COVID-19 specific listings from the DWMs in Table 3.2.</b> . . . . .	129

A.3	<b>Significant COVID-19 events.</b> We defined an acronym for each event and reported it in the main text plots. Please note that this list does not intend to be exhaustive or to establish a ranking between events. . . . .	132
A.4	<b>List of all DWMs, together with their specialization and a brief description.</b> . . . . .	135
A.5	<b>COVID-19 related medicines appearing in the listings,</b> together with a brief description and the number of listings related to that drug. . . . .	136
A.6	<b>COVID-19 related <i>guides on scamming</i> appearing in the listings,</b> together with a brief description and the number of listings related to that sub-category. . . . .	136
B.1	<b>Generic vaccine offered on a DWM.</b> Example of a listing in the <i>unspecified vaccines</i> category offering a generic vaccine, which does not specify the producer. Personal information of the vendor are hidden with # symbols. . .	138
B.2	<b>Translation of the proof of vaccination offered on a Hydra.</b> English translation of the listing in Fig B.2(b). We use Google Translator to translate the text from Russian to English. . . . .	139
B.3	<b>Example of drug listing mentioning COVID-19, and problems with the United States Postal Services (USPS).</b> . . . . .	140
B.4	<b>Example of drug listing mentioning COVID-19, and ensuring safety measures are taken.</b> . . . . .	141
B.5	<b>Example of drug listing mentioning COVID-19, and “stocks are almost exhausted by Corona”.</b> . . . . .	142
B.6	<b>List of DWMs analysed.</b> . . . . .	143
B.7	<b>Vaccine listings detected on DWMs.</b> Some vendors and DWMs offer vaccines that belong in more than one category. . . . .	144
B.8	<b>COVID-19 related products offered on DWMs.</b> Availability of COVID-19 related products since November 2020. . . . .	146
C.1	<b>Memory kernel fitted coefficients <math>\beta</math> and <math>c</math> for every degree class <math>k_{min} &lt; \text{degree} &lt; k_{min} * 2</math></b> . . . . .	163
C.2	<b>Details of the DWMs under study:</b> start, end, reason of closure, type, total number of transactions and total volume . . . . .	164
C.3	<b>Preferential attachment parameter <math>\Delta</math>.</b> Values of the preferential attachment parameter $\Delta$ for each product market, fitted with maximum likelihood estimation on the attractiveness distribution. . . . .	165
D.1	<b>Details of each attribute present in the S2AG dataset.</b> Number of papers (right column) in the <i>(de)centralization</i> dataset that contain information on the listed attributes (left column). The attribute <i>"inCitations"</i> and <i>"out-Citations"</i> correspond respectively to received citations and given references. . . . .	167

D.2	<b>Number of papers in decreasing order in the dataset that have been categorized with the respective tuple of fields of study (left table) and with the respective first field (right table).</b> Here we have listed only the fields of study with more than 1000 papers in the dataset. Tuples with more than two fields of study have been reduced considering only the first two. . . . .	168
D.3	<b>Results of the keywords annotation procedure.</b> Clusters at 3 <sup>rd</sup> level, with associated manually annotated keyword and number of publications, together with the same information for each of the clusters at hierarchical level 4 in its branch of the hierarchy. Single horizontal lines denote clusters at the 2 <sup>nd</sup> hierarchical level, while the double horizontal lines denote the three branches at the 1 <sup>st</sup> level. Clusters with less than 500 documents have been disregarded. . . . .	172
E.1	<b>General statistics of DWMs, part 1.</b> Some DWMs are presented here, the others are available in Table E.2. The terms “sent” and “received” always refer to transactions made by users. The trading volume indicates millions of dollars. The amount of dollars sent and received by users through U2U transactions is equivalent to the total. . . . .	179
E.2	<b>General statistics of DWMs, part 2.</b> Some DWMs are presented here, the others are available in Table E.1. The terms “sent” and “received” always refer to transactions made by users. The trading volume indicates millions of dollars. The amount of dollars sent and received by users through U2U transactions is equivalent to the total. . . . .	180
E.3	<b>Statistical tests.</b> The two-sided Kolmogorov-Smirnov test is used to perform the statistical test. All p-values are less than 0.0001, which is indicated with 0.0001. . . . .	189

# *Abstract*

Since the development of Bitcoin in 2008, blockchain-based technologies have flourished and several applications have been created, such as smart contracts, NFTs or cryptocurrencies. The blockchain allows for different kinds of transactions to be verified in a decentralized way, without the need for an intermediary, allowing new socio-technical ecosystems to be born and grow. Transactions are stored publicly, anonymously and openly on the blockchain, granting unprecedented access to data on human (collective) behaviour and socio-technical systems. The question thus arises: can we use this data to characterize these systems, and possibly to further our understanding of human behaviour? In this thesis, we address this question by studying different blockchain-based ecosystems through a combination of different datasets. Firstly, we study Dark Web Marketplaces (DWMs), online illicit markets on the dark web using cryptocurrencies for payments, and we characterize how they first reacted and then adapted to the COVID-19 pandemic using web scraping data. Secondly, we exploit a unique dataset of Bitcoin and proprietary transactions to characterize the buyer-seller network on DWMs and regulated e-commerce platforms. Thirdly, we study a comprehensive dataset of scientific publications to investigate the evolution of the concept of decentralization, pillar of blockchain-based ecosystems, in time. Then, we extend the literature on DWMs by studying the wider ecosystem of direct interactions between users, a network we can study only thanks to blockchain data. Finally, we analyse the trade of NFT collectibles on the largest open marketplace available, characterizing the role of rarity in determining market trends. Overall, this thesis presents a series of pioneering studies improving our understanding of blockchain-based socio-technical systems, thanks to unique comprehensive large scale datasets giving us unprecedented access to the history and behaviour of these ecosystems. We hope researchers will extend this work to improve our understanding of these systems and more generally human behaviour.

# Publications

This thesis is based on the following publications:

- 
- I. Alberto Bracci, Matthieu Nadini, Maxwell Aliapoulos, Damon McCoy, Ian Gray, Alexander Teytelboym, Angela Gallo, Andrea Baronchelli; *Dark Web Marketplaces and COVID-19: before the vaccine*; EPJ data science; 2021
  - II. Alberto Bracci, Matthieu Nadini, Maxwell Aliapoulos, Damon McCoy, Ian Gray, Alexander Teytelboym, Angela Gallo, Andrea Baronchelli; *Vaccines and more: the response of Dark Web Marketplaces to the ongoing COVID-19 pandemic*; PLOS One; 2022
  - III. Alberto Bracci, Jörn Boehnke, Abeer ElBahrawy, Nicola Perra, Alexander Teytelboym, Andrea Baronchelli; *Macroscopic properties of buyer-seller networks in online marketplaces*; PNAS Nexus; 2022
  - IV. Alberto Bracci, Gabriele Di Bona, Nicola Perra, Vito Latora, Andrea Baronchelli; *The decentralized evolution of decentralization across fields: from Governance to Blockchain*; Arxiv; 2022
  - V. Matthieu Nadini, Alberto Bracci, Abeer ElBahrawy, Philip Gradwell, Alexander Teytelboym, Andrea Baronchelli; *Emergence and structure of decentralised trade networks around dark web marketplaces*; Scientific reports; 2022
  - VI. Amin Mekacher, Alberto Bracci, Matthieu Nadini, Mauro Martino, Laura Alessandretti, Luca Maria Aiello, and Andrea Baronchelli; *Heterogeneous rarity patterns drive price dynamics in NFT collections*; Scientific Reports; 2022

Other publications:

- 
- VII. Giulia Martini, Alberto Bracci, Sejal Jaiswal, Matteo Corea, Lorenzo Riches, Jonathan Rivers, Elisa Omodei; *Machine learning can guide food security efforts when primary data are not available*; Nature Food; 2022
  - VIII. Benjamin Steinegger, Iacopo Iacopini, Andreia Sofia Teixeira, Alberto Bracci, Pau Casanova-Ferrer, Alberto Antonioni, Eugenio Valdano; *Non-selective distribution of infectious disease prevention may outperform risk-based targeting*; Nature Communications; 2022
  - IX. Alba Bernini, Elodie Blouzard, Alberto Bracci, Pau Casanova, Iacopo Iacopini, Benjamin Steinegger, Andreia Sofia Teixeira, Alberto Antonioni, Eugenio Valdano; *Evaluating the impact of PrEP on HIV and gonorrhoea on a networked population of female sex workers*; arXiv; 2019

# *Acknowledgements*

This thesis is the result of work I could have never accomplished alone. I want to thank:

Andrea Baronchelli, who has guided me throughout with years with care and has given me many opportunities and the freedom to pursue them.

My collaborators, whose experience and help have greatly shaped the research I have done, and who have taught me a lot. A special mention to Alex, Jorn, Nicola and especially Gabriele, who has shared these three years of PhD in London with me.

Prof. Mark Broom and Dr. Federico Botta, for kindly agreeing to be the examiners for this thesis.

Elisa and Giulia, and the wider WFP Hunger Monitoring Unit, for making me feel welcome at WFP from day one, and showing me that using data science for good is not a buzzword.

510, and in particular Corina and Iacopo, for giving me the chance to use my skills to help people, even remotely and during a pandemic.

Ed, Will, Taka, Alan, Michele and all the PhD students and postdocs in the department, for sharing a windowless office, good and tough times, and especially many lunches.

Amin, Maddalena, Elohim, Max, Laura and Matthieu, for sharing the joys and sorrows of our daily work, and especially many beers!

All my friends in Pisa and around Italy and Europe, for the thousand of texts, many calls and all the experiences we have shared.

Dhanya, for always being there during these years, always supporting me and helping me in so many ways it is impossible to list.

La mia famiglia, senza il vostro supporto non avrei fatto niente. Se sono qua adesso, lo devo solo a voi. Grazie.

# Chapter 1

## Introduction

Blockchain-based ecosystems are unique examples of decentralized socio-technical systems. The blockchain is a decentralized ledger where transactions are stored and verified in a secure, pseudo-anonymous way without the use of any intermediary. The most famous and first example is that of cryptocurrencies, which allow two people to exchange digital money without the use of a bank or any other third party involved. This new technology has rapidly evolved and given birth to many different applications and new social systems, where these decentralized interactions allow for the emergence of interesting and non trivial collective behaviour. For instance, the strong and volatile increase in trading value of many cryptocurrencies, among which Bitcoin, is the most clear example of social coordination, where people collectively agree on assigning value to a currency without any backing or strong power behind it. Importantly, the blockchain is by design public and pseudo-anonymous, in almost all cases, with data stored publicly and transactions available to all researchers. This allows for an unprecedented level of detail on new, rapidly evolving socio-technical systems, as well as more traditional ones, and opens up new avenues of research which have only been scratched so far.

The blockchain is a technology which first saw the light with Bitcoin in 2008 [3], the first example of cryptocurrency. Since then, thousands of cryptocurrencies have been created, with varying success. While Bitcoin has seen its value dramatically increase, even while being very volatile, many cryptocurrencies have tried to emulate its success. Among many, one that deserves a special mention certainly is Ethereum [4], whose differences with Bitcoin allow it to be more versatile and able to spawn many more applications. In particular, its ability to execute smart contracts has allowed the birth of several new markets: gaming, Non Fungible Tokens (NFTs) and decentralized finance (DeFi) being among the most famous and largest markets. For instance, Bitcoin reached a price of 61k USD in 2021 [5], DeFi had a market volume of 120B USD in February 2022 [6]. NFTs on the other hand represent a market reaching a monthly trading value of 5B USD in January 2022 [7]. Due to the absence of intermediaries and its anonymity features, Bitcoin, as well as other cryptocurrencies, is also used for illicit activities. For example, in addition to numerous scams, Bitcoin is the

main currency on Dark Web Marketplaces (DWMs), online marketplaces for illicit goods, a market worth over 4B USD up to 2019 [8].

The literature on blockchain based ecosystems is still young, and naturally focuses on the aspects which captured the public attention in their striking rise. For instance, the largest majority of the literature either focuses on the details of the underlying technology, such as which proof of x algorithm to use [9, 10], or on understanding the reasons and mechanisms behind the different market trends and price evolution [11, 12, 13]. Overall, cryptocurrencies have attracted the largest fraction of the related research due to their popularity. Yet, other blockchain based ecosystems such as Dark Web Marketplaces (DWMs) [14, 15, 16] and NFT markets [17, 18, 19] have stimulated research in recent years. However, few articles have employed a complex systems approach trying to understand the collective behaviour of these systems and how they interact with other systems and society at large, and many questions remain unanswered, and this has also been stimulated by recent events such as the COVID-19 pandemic.

In this thesis, we build on the recent yet growing literature on blockchain-based ecosystems, extending previous results and opening new lines of research which other researchers will hopefully build upon. Our contributions can be summarised in four main points encompassing the main chapters of this thesis. Firstly, *we improve our understanding of multiple blockchain-based ecosystems*, including Dark Web Marketplaces, their interplay with economic shocks and their behaviour compared to regulated marketplaces; the user to user network growing and coexisting around DWMs; NFTs and their market trends; and the wider scientific literature at the foundations of the blockchain itself. Secondly, *our analysis covers an extensive period of time, including major events for society and the system itself*: the full history of DWMs from 2011 to 2021, including their creation, evolution and reaction to the pandemic, with the ability to compare with one regulated online marketplace across the same period of time, but also the major market boom of NFTs from 2018 to 2022. Thirdly, *we use a complex systems approach to be able to investigate the collective behaviour of these complex socio-technical systems*, using tools from network science, probability theory and statistical physics, and proposing novel models and methodologies combining tools from all these areas to analyse and explain the behaviour of such systems. Finally, *we collect and pre-process unique comprehensive datasets across multiple systems, from both commercial companies and openly available datasets*, including listings and transactions from more than 100 DWMs across over 10 years, purchases from one regulated e-commerce platform over 10 years and multiple product markets, the full scientific literature on decentralization since the 1950s, direct transactions between users of DWMs since their inception, and more than 400 NFT collections transactions between 2018 and 2022,

Our research addresses questions of interest around different blockchain-based ecosystems, how they self-organize and their interplay with wider society, while also addressing themes of more theoretical and methodological interest. The research is based on novel unique datasets of unprecedented size and coverage of the studied systems. For instance, chapters 3 and 4 are based on a dataset containing listings scraped from more than 100 Dark

Web Marketplaces since 2020, providing the best coverage available of the whole ecosystem with a daily resolution, overcoming several limitations of past studies and allowing us to study the whole ecosystem at once. Chapter 5 is instead based on two datasets of transactions covering the 28 most important DWMs of the whole ecosystem history and 144 product market of one large regulated e-commerce platform respectively, covering the period from 2010 to 2021 (i.e. the full history of the DWMs ecosystem). An extension of the DWM dataset is also used in chapter 7, where we include also transaction between users of the markets, covering a larger set of markets including minor platforms, and was only possible thanks to the nature of the blockchain (i.e. recording all transactions ever done). This data was accessed thanks to agreements with private companies, due to the difficulties which make gathering and processing such data by a research group impossible, and represents a unique opportunity to investigate the collective behaviour of these blockchain based ecosystems on such a scale and resolution. However, this thesis also uses openly collected and available datasets. For instance, chapter 6 exploits the openly available S2AG dataset [20], which aims at indexing the whole scientific literature, allowing us to study a collection of almost 200 thousand publications mentioning the concept of decentralization, requiring us to develop an ad-hoc pipeline (publicly available <sup>a</sup>) to analyse such a large collection of documents. Finally, chapter 8 is based on a dataset covering 400 NFT collections between 2018 and 2022, collected by us for the scope of this study, and forming the largest openly available dataset on the topic of collectible NFTs, which we made available to all interested researchers through an open repository <sup>b</sup>. More details on each dataset are present in each chapter and relevant appendix.

To analyse such datasets, this thesis has relied on, and improved upon, state of the art techniques coming from statistics, applied maths and physics. For instance, chapters 3 and 4 employ different statistical techniques and methods from Natural Language Processing [21], such as sentence embeddings [22] and dimensionality reduction [23] techniques, to analyse the large volume of listings data available from almost 200 DWMs. In chapter 5 the main framework utilized to conduct both the data analysis and mathematical modelling is that of complex networks [24], in particular by using bipartite networks and activity driven network models [25] to model transactions on online marketplaces, and to simulate their dynamics and the behaviour of their users. In chapter 6 we again used a combination of statistics and complex networks to model the citation dynamics [26] and flows of knowledge [27] between different fields of the scientific literature. The framework of complex networks has also allowed us to understand the dynamics of user to user interactions around DWMs in chapter 7, in particular by using temporal network models [1, 28], while statistical analysis has instead been the main tool in chapter 8.

Here, we present how the thesis is structured around five main questions, to which each chapter represents a (partial) answer. In chapter 2, an overview of what blockchain technology is, its main applications relevant to this thesis, and the relevant literature are presented.

---

<sup>a</sup><https://github.com/alberto-bracci/decentralization>

<sup>b</sup>The data can be downloaded at <https://osf.io/7w9r6/>

In chapters A-F, appendices to each main chapter are presented, containing additional material, robustness tests, and additional figures and tables.

*Chapter 3 and 4: How Dark Web Marketplaces reacted, and adapted, to the COVID-19 pandemic?*

Dark Web Marketplaces (DWMs) have been shown to be a main channel for illicit trade, exploiting capabilities of anonymous browsing (TOR) and cryptocurrencies anonymity. Their trade has steadily increased in time, and new markets have constantly emerged to replace those closing because of scams or police raids [29]. In this context, a natural question to ask is how they adapted to the COVID-19 pandemic, a shock which enormously impacted the regulated economy affecting the supply chain of several goods, creating barriers between countries and regions, and most importantly creating the demand for specific products. We address this question by analysing a unique dataset of listings information scraped directly from multiple DWMs by Flashpoint Intel [30]. We analyse 851,199 listings extracted from 30 DWMs between January 1, 2020 and November 16, 2020, including all major markets at the time. We identify 788 listings directly related to COVID-19 products and monitor the temporal evolution of product categories including Personal Protective Equipment (PPE), medicines (e.g., Hydroxycyclochlorine), and medical frauds. Then, we show how temporal trends in their availability and public attention correlate, as measured by Twitter posts and Wikipedia page visits. Finally, we investigate the impact of COVID-19 on other listings, showing how pandemic related events led to an increase in COVID-19 mentions on DWMs. After showing how DWMs reacted to the emergence of the global COVID-19 pandemic, it is interesting to study how they kept adapting to it in the following months. We address this by studying an extended version of the first dataset, containing listings from 194 DWMs collected until July 2021. We first show how DWMs adapted to the changing landscape of the pandemic by offering officially approved vaccines listings, like Pfizer/BionTech, as well as COVID-19 passports and fabricated proofs of vaccinations. We then show how these products have replaced other COVID-19 related products initially offered on these DWMs, pointing again to a clear link between shortages, public attention and products offered on DWMs. Finally, we investigate in more detail the impact on non COVID-19 related listings, showing how drugs have been the most affected goods.

The content of this chapter is based on publications [I] and [II].

*Chapter 5: What is the structure of buyer-seller networks on online marketplaces, and does it differ between unregulated (blockchain-based Dark Web Marketplaces) and regulated markets?*

Online marketplaces, like Amazon or DWMs, are the main platforms for legal and illegal e-commerce. We have seen how DWMs react and adapt to sudden economic shocks, however, their general empirical properties, and those of regulated online marketplace, are poorly studied due to the lack of comprehensive large scale transaction datasets. Here, we exploit two unique datasets containing a total of 245M transactions (total volume of 16B USD) from both regulated and unregulated platforms between 2010 and 2021, including 28 DWMs

and 144 markets from a regulated e-commerce platform. This data allows us to investigate the main properties of buyer-seller networks, a useful tool to study the behaviour of users on these markets, and more in general the behaviour of online marketplaces across different platforms and over an extended period of time. First, we show that the main macroscopic properties of buyer-seller networks forming on these platforms are regular across different instances of the same kind of platform, but also across regulated and unregulated platforms, pointing to a possible universality ultimately due to human behaviour. We then investigate possible mechanisms behind such properties, uncovering the role of memory in driving buyer behaviour. Finally, we exploit these observations to build a mathematical network formation model which reproduces the main properties of the empirical networks, showing the fundamental role of memory and preferential attachment mechanisms.

The content of this chapter is based on publication [III].

*Chapter 6: What is the origin and evolution of the concept of decentralization across different academic fields?*

Decentralization is one of the pillars of blockchain technology, promising a world where transactions of any kind can be executed directly between two parties without the need for a third party to verify them. However, it is often unclear what decentralization actually means, and this concept is actually widespread across different disciplines such as Economics, Political Science and Computer Science with different meanings. Here, we try to shed light on the origin and evolution of this concept by studying a dataset comprising of 425k academic publications mentioning the concept. First, we uncover that the fraction of papers on the topic has been exponentially increasing in the past 70 years, reaching 1 author in 154 in 2021. We then employ a network based clustering method, using both semantic and citation information, to hierarchically cluster papers and topics, showing how the concept of decentralization has independently emerged in different fields, with cross-contamination emerging only in recent years. Finally, we employ the methodology of knowledge flows [27] to study how information has flowed between different fields, showing how Blockchain has become the most influential field in the past 10 years, while Governance played a dominant role until the 1990s. Importantly, we publicly released the code of the pipeline used for the analysis performed in this chapter, allowing other researchers to use it to study potentially any other concept.

The content of this chapter is based on publication [IV].

*Chapter 7: Do Dark Web Marketplaces user trade directly among each other, without using the market platform?*

We have already mentioned the central role of DWMs in the online trade of illicit goods, we have investigated the macroscopic properties of the buyer-seller networks forming on these platforms, and we have shown their importance and how they swiftly reacted and adapted to the COVID-19 pandemic. A natural question, which generalizes to online marketplaces, is whether buyers and sellers who meet on these markets start to trade directly between each other, bypassing the market role as intermediary. In other words: is the market a meeting

place for users, who then transact directly between each other? Here, we extend previous interview-based studies by studying a dataset of 31 million Bitcoin transactions among users of 40 DWMs between June 2011 and Jan 2021, collected by Chainalysis [31]. First, we find that half of the DWM users engage in direct transactions without the intermediary role of the market, generating a volume larger than the DWM ecosystem itself. We then show that a relevant fraction of them form stable trading relationships, and generally come from the segment of DWM users with larger trading volumes. We also show how these trading pairs often form after both users have traded on the market. Finally, we demonstrate how the pairs keep engaging after the DWM closure, suggesting the presence of a more resilient decentralized network of trade around DWMs.

The content of this chapter is based on publication [V].

*Chapter 8: What is the market dynamics among collectible NFTs?*

The previous chapters have mainly studied Bitcoin based systems, yet 2021 has seen the boom of a new blockchain-based ecosystem and market based on Ethereum: NFTs. NFTs allow for digital object to be uniquely certified on the blockchain, and create new mechanisms for artists to get their work recognised. We investigate a relevant subset of the market: collectible NFTs, usually grouped in collections sharing common traits and features, and often algorithmically generated. Anecdotally, this market has seen some collection or specific NFTs gain incredible popularity and prices, yet the reason why it happens escapes traditional art market logic. Here, we address this by studying the role that rarity plays in the market dynamics, by studying a unique dataset of 3.7M transactions collected between January 2018 and June 2022, involving 1.4M NFTs distributed across 410 collections. We quantify rarity based on human-readable attributes, and show how it's generally heterogeneously distributed within a collection. We then analyse the market performance of NFTs, demonstrating how rare NFTs on average sell for higher prices, are traded less frequently and get higher returns on investment with less associated risk.

The content of this chapter is based on publication [VI].

Overall, our research improves our understanding of blockchain-based ecosystems by answering these questions. However, these results also call for more research to investigate new questions that stem from it. Future work can move along multiple directions. For instance, more research can be done to investigate the behaviour and dynamics of the emergent user to user network around DWMs, and whether similar networks also exist around regulated online marketplaces. This and other questions are explored in chapter 9, where we recap the main contributions of this thesis and detail some possible lines of research which could further improve upon it.

# Chapter 2

## Background

While only formally introduced by the Bitcoin white paper in 2008 [3], the young scientific literature on blockchain and its applications is rich and growing. Yet, as previously mentioned, it tends to be skewed towards certain topics, with a clear preference for the financial aspects of it driven by the volatile and high price of cryptocurrencies. While the aim of this thesis is to improve our understanding of the behaviour of complex blockchain-based socio-technical systems, one cannot do so without first delving into what the blockchain is, its recent history and main applications, and the existing state of the art literature on the systems studied in this thesis. In this chapter, we will endeavour to do so by taking an historical point of view. First, we will review Bitcoin, the first cryptocurrency, what the blockchain actually is and how it works in Bitcoin. Then, we will look at the second most famous cryptocurrency: Ethereum, and how it differs from Bitcoin, setting up the stage for modern blockchain applications such as NFTs and the metaverse, which we will briefly describe. A section will be dedicated to illicit uses of blockchain technologies, with particular care dedicated to Dark Web Marketplaces, a system which we will study in great detail throughout the thesis. Finally, we will dedicate a section to describing the main data sources which are used to study the blockchain based ecosystems which are the subject of this thesis, going over their main properties, what they allow us to study, and their main limitations.

### 2.1 Where it all began: Bitcoin

In 2008 a white paper titled "Bitcoin" was circulated on the mailing list Cyberpunk, authored by an unknown author named Satoshi Nakamoto, followed in the same year by an open source implementation of the ideas described in the same paper [3]. The project proposed a decentralized secure digital currency, based on technological innovations in databases, cryptography and network protocols [32]. Overall, Bitcoin is the result of three different technologies: the Blockchain [33], hash functions [34] and peer to peer networks [35].

Bitcoin is a digital currency that operates in the absence of banks or any intermediary.

However, it still needs to record all transactions in a secure ledger. The Blockchain is the secure ledger where all transactions ever made are recorded. The Blockchain is an actual chain of blocks, where one block records a number of transactions and a pointer to the previous block in the chain. As a decentralized system, the Blockchain is stored on all computers participating in the Bitcoin network; that is they are running the Bitcoin Core open-source software. This guarantees the safety of the data against failures and possible cyberattacks.

Even if all transactions are open and public, cryptography guarantees the anonymity of the users, since an address is stored on the blockchain instead of names or any personal information. The address is a collection of letters and numbers. Each user can potentially create a new address for every transaction, linking them to their wallet. No personal information is either linked to the wallet or needed to perform transactions or run the Bitcoin software. When a user makes a transaction, users participating in the network are automatically informed of it, and the transaction is put into a pool of unverified transactions, which need to be verified and then written into the blockchain.

Transactions are verified in blocks, which are added to the blockchain. To ensure that tampering with the blockchain is hard, if not impossible, miners (users of the blockchain peer to peer network verifying transactions) need to solve a computationally hard puzzle when verifying a block of transactions, which essentially boils down to guessing random numbers as quickly as possible. The solution of the puzzle is the output of a SHA-256 hash function run on the block and the pointer to the previous block. The function is hard to compute, but it is easy to verify whether the solution is correct. This "guessing" process on average takes 10 minutes, and its difficulty is periodically adjusted to account for periodical increases in computational power and hardware improvements. This way, if one were to tamper with an older block (e.g. change a transaction amount), they should also recompute the hashes for all subsequent blocks, taking a lot of time and power. Moreover, they would need to convince more than half of the network to accept their new chain in place of the correct one, making tampering with the blockchain essentially impossible in theory. Obviously, if one had control of more than half of the network, that would be possible, but at the same time trust in Bitcoin would fall together with its value and use (an argument of this kind was originally made in the Bitcoin whitepaper [3]). To incentivize mining, which we remark basically replaces the work of banks and intermediaries, miners are rewarded for their work with some Bitcoins, whose quantity is halved every 210,000 blocks to keep the final quantity finite and control the inflation rate. Miners can also gain Bitcoins from fees, which are however not compulsory. This whole mechanism is known as Proof of Work. As we'll see later, many cryptocurrencies share the protocol and principles of the blockchain, yet not all of them implement the proof of work mechanism. Other currencies, such as Ethereum will soon do, use proof of stake or other alternative mechanisms.

## 2.2 Dark Web Marketplaces

The previous section has described Bitcoin as it is described in the white paper and implemented as a technology. However, Bitcoin is not simply a technology, as it also involved economic incentives and human behaviour to ensure its functioning as described and designed. In reality, Bitcoin has evolved far from being a widely adopted digital cash, and has had many hiccups along its way. For instance, its claim of decentralization [36] and anonymity [37] have been constantly challenged, and its limited transaction speed, massive electricity usage and complex technology have so far stayed in the way of massive adoption. Other issues that research has highlighted include the central role of core developers in determining the future of the currency [38], the unequal distribution of wealth [39] and the limit to anonymity due to the ability to cluster addresses [40] among others. While it is out of the scope of this thesis to take a deep dive into all these issues, their causes and details, it is important to denote an area that has predominantly adopted Bitcoin as its main currency: illicit activities, and in particular Dark Web Marketplaces (DWMs) [41, 42], which we will study in detail throughout the thesis.

The online shadow economy is as old as the Internet. The first reported illegal online deal involved drugs and took place in 1972 [43]. The World Wide Web [44] facilitated the emergence of online illicit marketplaces [45, 46] but the first marketplaces could not guarantee anonymity and therefore permitted the traceability of users by law enforcement [47]. Modern DWMs originated and still operate online, but outside the World Wide Web in an encrypted part of the Internet whose contents are often not indexed by standard web search-engines [48]. *Silk Road* marketplace, which launched in 2011, was the first modern DWM [49]. It proposed a new way of trading drugs and other illegal products online and anonymously [50, 51, 52]. There were two key ingredients of *Silk Road*'s success. First, potential customers could access it using the Tor browser [53], which made their traceability difficult. Second, and this was the real innovation that made it possible, purchases were made in Bitcoin [3], which provided to the transactions a degree of anonymity which was previously simply unavailable [54, 55, 56].

After the FBI shut down *Silk Road* in 2013 [57], new DWMs quickly appeared (and later closed voluntarily or because of police operations<sup>a</sup>), offering drugs, firearms, credit cards, and fake IDs [58]. These DWMs also adapted to further increase the level of privacy and security offered to users [59, 60], such as the Invisible Internet Project (I2P) [61] and escrow checkout services [62]. Tor, now available for mobile devices as well, still offers more privacy than many other popular mobile applications [63] and Bitcoin is currently the most popular cryptocurrency in DWMs [41, 42, 64]. Trade today on DWMs is worth at least several hundreds of millions USD per year, and involves hundreds of thousands of buyers and vendors [57, 65, 8, 66, 67, 68]. As a result, law enforcement agencies have put considerable effort into combating them [57, 67, 68]. Furthermore, DWMs have been targets of cybercriminal actors through use of distributed denial-of-service (DDoS) attacks,

---

<sup>a</sup>see [https://www.emcdda.europa.eu/system/files/publications/8347/Darknet2018\\_posterFINAL.pdf](https://www.emcdda.europa.eu/system/files/publications/8347/Darknet2018_posterFINAL.pdf) for a complete history of the first years of the DWMs ecosystem

hacking attempts, and some DWMs also shut down due to administrators stealing funds from customers directly [69, 70]. However, DWMs have organised into a robust ecosystem which has proven exceptionally resilient to closures and whenever a DWM is closed, the users trading higher volumes of Bitcoins migrate to active DWMs or establish new ones [8].

The resilience and functioning operations of modern DWMs are possible partially because of numerous websites and forums where users can share their experiences. One example is Dread [71], a Reddit-like forum created in 2018 after the closure of the dedicated pages on Reddit [72]. Other ad-hoc platforms exist to monitor whether known DWMs are active or currently unavailable [73, 74, 75, 76]. Additional mechanisms, like feedback and ratings systems [65], enhance the resilience of these DWMs and build trust towards the DWM and its vendors. In a similar way to regulated online marketplaces, DWM buyers are asked to leave feedback and a rating after a purchase. Additionally, DWM administrators often act as vendor moderators by banning vendors or specific categories of products.

DWMs have been used to circumvent laws and regulations. They have been the subject of many case studies [59, 60, 77] and comparative analyses [50, 51, 52, 78, 79, 80]. These studies highlighted that illicit online transactions in DWMs are perceived as safer than negotiating in face-to-face drug markets [59]. They are based on the concept of “harm reduction,” where vendors prefer to sell tested and high quality products [60]. Vendors customize their products to match the specialisation of different DWMs thus creating an efficient distribution network [80], which sometimes goes beyond a base retail market [79]. While these characteristics favour the DWM economy against the offline shadow economy, DWMs sell a variety of illicit products [50, 51, 52, 78], such as, drugs, fake IDs, “how to” manuals (for scams, bombs etc.), and weapons. One prominent category is that of digital goods [81], including ransomware, social engineering guides, and financial malware to steal credit cards and bank account credentials. While research has improved our understanding of DWMs under many aspects, the lack of extensive and updated datasets has left some open questions such as: how do markets react and adapt their offer to sudden external economic shocks? What are the main economic empirical properties of these platforms? Do these platforms foster a direct user to user trade that bypasses the markets themselves?

## 2.3 Ethereum

As previously mentioned, Bitcoin paved the way for thousands of cryptocurrencies to be created with varying success. Among them, the most innovative and successful certainly is Ether, whose blockchain and underlying technology is called Ethereum. Ethereum was released in 2015 [4], and represents the most important technological innovation since Bitcoin. While it’s originally based on the same proof of work mechanism (changing to proof of stake at the time of publication of this work), with some technical differences such as the hashing algorithm or the block waiting time, and can work as a medium of exchange exactly as Bitcoin, it allows for more functionalities to be built on top of it. Indeed, Ethereum is a programmable blockchain, meaning that it has a fully functioning programming language

which allows code (and obviously data) to be stored and executed in a decentralized way on the blockchain, theoretically with no limits to what it can do.

Programs running on the blockchain are called smart contracts. They are executed when triggered by a transaction from a user (or another smart contract). Once published on Ethereum, they cannot be removed, and they will be operational for as long as programmed or Ethereum exists. Not even the original programmer can remove them, or stop them from executing their code, as this is a core founding feature of the blockchain (exactly like Bitcoin transactions cannot be tampered with). This feature has caused issues in the past. One of the most famous example is that of the DAO hack [82]. The DAO was a first attempt at a Decentralized Autonomous Organization, an organization collectively owned by its members, whose rules are set and automatically executed on the blockchain. The DAO was built as a venture fund, and by 3 weeks of its token sale it had raised over 150M dollars. However, a bug in the code of the smart contract made it possible for a hacker to steal the funds. This event caused a huge debate in the community, seeing people advocating for the funds to be recovered (by reverting the blockchain state to before the hack), and others claiming that "the code is law", with the hack being legit as it only did what the contract made possible. The event eventually led to a hard fork, where Ethereum reverted to its previous state, while a fraction of its users continued the state of the blockchain containing the hack, that came to be known as Ethereum Classic [82].

Since then, the Ethereum ecosystem has flourished and evolved, with numerous applications developed in multiple fields. One of the most successful ones certainly is Decentralized Finance (DeFi), which is also advertised as one of the main applications of Ethereum by the Ethereum Foundation [83]. DeFi aims to constitute an alternative to traditional finance, giving access to financial markets without the presence of any intermediaries, in the spirit of the blockchain. Through DeFi, people gain access to traditional financial products like mortgages and trading, without banks or governments possibly denying access. DeFi has gained traction in recent years, reaching a market cap of 120B USD in February 2022. Another major application of Ethereum is that of Non Fungible Tokens (NFTs), a market which gained public attention after the artist known as Beeple sold an NFT for 69M USD in March 2021 [84]. NFTs are digital assets that represent objects like art, encoded in smart contracts, certifying the uniqueness of the digital asset and therefore also the ownership of it. NFTs represent a new avenue for artists, who can exploit the potentialities of the blockchain and smart contracts to gain value and recognition for their digital pieces. For instance, smart contracts can be used to enforce that a fixed share of each secondary sale goes back to the original artist. While NFTs have encountered some difficulties, as many application of such a young technology as the blockchain, their fast rise has made them a strong reality, reaching a volume of 5B USD of sales in January 2022 [7].

Among the multiple dimensions and applications of Ethereum, chapter 8 will present some results around the dynamics of the NFT market. Research on NFT is understandably young, given the recent rise of the phenomenon, yet it is growing rapidly mirroring the growth of NFTs. Previous studies include an overview of the overall market, trade networks, and

visual features of NFTs, and their impact on price prediction [17], showing for instance the strong heterogeneity of the market, with the top 10% of the traders performing 85% of the transactions. Other research has focused generally on specific marketplaces or collections. For example, Vasan et al. [18] have study the Foundation marketplace, uncovering how artists invited on the platform by other successful artists are also successful, and how a small cluster of collectors repeatedly invests in such artists. The role of social media attention has also been investigated [85], showing how the artists' presence on twitter can be a signal of NFT success. Other studies have instead looked at the creators-collectors network [19], and the financial advantage of experienced users [86]. Along this line, research also suggests that NFTs have become a promising investment as a fintech asset [87]. Other lines of research include the analysis of illicit transactions connected to NFT trading [88, 89] and of their connections with financial indicators [90, 91, 92]. The metaverse, an NFT submarket which has recently garnered attention both from big tech companies [93] and popular NFT creators [94], is another focus of research [95, 96]. However, what determines the success of an NFT, in contrast to other less successful ones, is still an open question.

## 2.4 Data sources

In this section we describe the main data sources which have been used, and we have used, to study the blockchain-based ecosystems subject of this thesis. The goal of this section is to provide a detailed enough background of the main data sources, without going into the specific numbers and details of the datasets used in each study, which differ from chapter to chapter and will be sufficiently described in each of them. After this section, the advantages and limitations of each kind of data will be clear enough, making it easier to understand the context of each chapter. Data not strictly related to blockchain-based ecosystems, such as the regulated online e-commerce marketplace transactions used in chapter 5 or the scientific publications used in chapter 6, will not be discussed in this section, as their understanding is generally more straightforward and their exposition should be clear from the description in the relative chapters.

### 2.4.1 DWM listings

The first and most popular data source used to study DWMs are DWM listings. A listing consists of the advertisement for the sale of a product on the DWM platform, containing information such as the title, product description, advertised price, seller username and other potential attributes, such as shipping information and geography, category of the product or number of available items. Listings information can be scraped from the DWMs' websites. Listings data therefore consists of tabular data containing information from the listings directly scraped from the DWMs pages. This was done extensively in the early years of the platforms and related research, but has been made difficult by the DWMs in recent years, requiring dedicated tools and effort to circumvent their protection measures. An example of listing, highlighting useful information which can be scraped from it, is shown in Fig. 2.1.




Figure 2.1: **Example of a DWM listing.** Screenshot of a chloroquine listing in the DarkBay/Dbay marketplace, where we highlight some of its salient attributes. Among the attributes considered in this work and shown in Table A.1, “Time” and “Marketplace name” attributes are not present in this screenshot, while the “Quantity” attribute is not fixed by the vendor.

Listings data has allowed extensive research on DWMs, giving a major contribution to their current understanding. For instance, listings information has allowed researchers to characterize the DWM landscape in specific countries in terms of vendors and products/prices [97, 98]. Other work has instead focused on understanding the effect of market closures and law enforcement operations on the DWM ecosystem. These markets have been the target of police operations since the beginning [57, 99, 100], while some markets have either simply shut down or suddenly closed by stealing the users’ money which were deposited into their accounts (a phenomenon usually known as exit scamming). In this context, new markets have constantly risen to fill the gap [101], and research based on listings data has shown how the ecosystem has constantly grown, and "the effect of law enforcement take-downs is mixed at best" [102]. Not surprisingly, a significant number of studies has focused on exploiting the detailed information available from listings to characterize the kinds and trade of goods available on these platforms. Illicit drugs have been found to be the major product on multiple platforms [103, 104, 105], with cybercrime products [106], forged identity documents [103] and credit card dumps [102], while other studies have shown the strong impact of geography constraints [107, 108], showing for instance how DWMs are likely to be a vector for sale between dealers and consumers, not facilitating a direct sale between

producer and consumer. In this thesis, namely in chapters 3 and 4, we will use an updated and large scale dataset of DWM listings, covering almost 200 markets, to study how the COVID-19 pandemic has impacted the offer of goods on DWMs.

However, listings data present several constraints which have limited the studies using such data source, and have impacted their conclusions. For instance, listings crawled data rely on what is declared by the seller and available on the market platform. This is a strong limitation when dealing for example with geography data, but also with price. Indeed, it has been found that sellers often arbitrarily increase the price of a product when their stock is not available, in order not to take the listing down and lose the product ratings [103]. On the opposite end of the spectrum, vendors may take down listings with low ratings to just advertise new listings for the same product, making it hard to do longitudinal analysis on a specific good. Other limitations include the difficulty of overcoming the platform activity to stop automatic web scraping, often leading to datasets focusing on short periods of time and few platforms, or presenting significant gaps in the data. Finally, one major limitation of the previously mentioned work relates to the use of listings to estimate sales. Multiple works [103] have used listing reviews as a proxy for sales, assuming that one review correspond to the sale of one unit of product. Such assumption is a lower bound at best, but due to limitations of reviews, such as not all users leaving one after a sale, or users bulk buying, still leads to estimates whose error is hard to quantify. These limitations can be largely overcome if one were to have access to the raw transaction data relative to this market activity. While having access to such data from seized servers would require collaboration with law enforcement agencies, and would give access only a posteriori to closed DWMs, one can exploit the fact that these platforms use Bitcoin as their main currency. In the next section, we will detail how Bitcoin transactions can, and have been, used to study DWMs on a larger scale and finer details.

## 2.4.2 Bitcoin transactions

The raw, anonymised Bitcoin blockchain can be publicly accessed through Bitcoin core [109] or third-party APIs such as Blockchain.com [110]. It contains information about origin and destination addresses, as well as time and amount of the transactions. In order to contrast traceability of the real identity, a user is likely to use multiple addresses. A new address is often generated in each transaction. Grouping the addresses in clusters reduces the complexity of the Bitcoin blockchain and challenge users' anonymity [111]. Given that millions of Bitcoin addresses are currently active and many others are continuously being generated, a clustering approach primarily based on manual annotation is not feasible. Various heuristics, instead, have been proposed [111, 40, 112, 113]. They were successful in grouping Bitcoin addresses and associate them to cluster of real entities. For instance, in [111], the authors were able to find a connection between a set of large transactions and a single one, which was dated in November 2010. In [40], the authors applied to a daily university setting the privacy protocol recommended in Bitcoin transactions, finding that almost 40% of the real identities would be recovered. Another work showed the presence of "super clusters" of entities, which

marked macro-variations in the evolution of the Bitcoin economy [112]. The primary reasons behind the effectiveness of heuristic clustering are: “address reuse, avoidable merging, super-clusters with high centrality, and the incremental growth of address clusters” [113].




Figure 2.2: **Identification of real entities in the Blockchain.** End goal of Bitcoin transactions clustering techniques: mapping a series of Bitcoin addresses to real entities. In this example, an address sends Bitcoins to another address. Thanks to the identification process, the two addresses are associated with two real entities. The Bitcoin transaction between the two entities becomes traceable and transparent.

The end goal of clustering Bitcoin addresses is to map them to single, real entities, as shown in Figure 2.2. To achieve this goal, however, heuristic clustering techniques should be improved. Manual annotation has shown a valuable potential [114]. It consists on gathering publicly available Bitcoin addresses, like the Wikimedia Foundation one [115], and engage through direct interaction with unknown Bitcoin addresses. If some real entities are known, it is easier to associate the remaining Bitcoin addresses to other real identities. In the last few years, companies specialising in Bitcoin analytics have started to leverage previous methodologies [111, 40, 112, 113, 114] to unveil real entities. The leading company in analysing Bitcoin transactions on DWMs is Chainalysis Inc. [31], which has also aided several federal investigations. For instance, it supported the United States Internal Revenue Service (IRS) in tracking Bitcoin transactions [116] and the FBI in the Twitter hack [117]. Chainalysis clusters Bitcoin transactions in groups by combining previous methodologies [111, 40, 112, 113] and real entities are unveiled with an approach similar to [114].

In our work, we rely on data collected and pre-processed by Chainalysis Inc. [31], a company specialising in blockchain analytics, where they cluster addresses into entities, and identify entities corresponding to DWMs. The dataset, or rather a less updated version of it, has already been used in previous work [8]. The data is collected using a complex network perspective. As we’re interested in studying DWMs, we start from those entities labelled as DWMs. All transactions involving these entities are collected, which identify the nearest neighbours, or users, of the market. Transactions involving exchanges are not included, to limit the size of the dataset and analyse only user behaviour. Then, all the transaction history of the nearest neighbours is also collected, thus including next nearest

neighbours but only when they interact with nearest neighbours. Each market network is therefore an egocentric network of radius 2, and the whole dataset is a collection of these smaller ego-centric networks.

Such data allows for a more complete analysis, or rather complementary, than what listings allow. Indeed, one has more reliable measures of the buyer activity, sales and the overall market volume, without access to information on the single purchases. When combined with other metadata, such as market openings and closures, it allows for a complete picture of the market history. For instance, [8] has provided the first complete picture of the DWM ecosystem from SilkRoad, the first market in 2011, to 2019. It quantitatively showed how the market activity had been resilient to the market closures, both in terms of volume and users, with active users migrating to other existing DWMs when a market closed. In this thesis, namely in chapters 5 and 7, we use updated versions of the same dataset to further explore the behaviour of these platforms and its users.

However, while overcoming limitations of listings data, this data source presents some limitations as well. For instance, no ground truth is publicly available to benchmark and evaluate the clustering method. While the methods used by research and companies are state of the art, they still present limitations which hinder the clustering of addresses. For example, using this clustered data in police investigations [118] means that it is crucial to avoid false positives (i.e. cluster together two addresses which should not be), meaning that on the opposite side sometimes addresses are not clustered together even if they should have.

### 2.4.3 NFT transactions

As previously mentioned, Ethereum transactions, as well as those of other blockchains modelled on Ethereum, allow for more functionalities than the simple exchange of money or value. The main application of this that we will study in this thesis is that of NFTs. As previously mentioned, "an NFT is a unit of data stored on a blockchain that certifies a digital asset to be unique and therefore not interchangeable, while offering a unique digital certificate of ownership for the NFT" [17]. NFTs have been prominently used in the art market, but also in gaming and in the fast growing world of the metaverse. Due to their nature, any transaction related to NFTs is therefore recorded on the blockchain, and therefore can be publicly accessed and studied. However, due to how transactions and in particular NFT metadata are stored on the blockchain, it is not easy to massively download all NFT related transactions directly from the publicly available blockchain data. Yet, APIs are made available by NFT marketplaces, or sometimes websites gathering aggregated information on the NFT world make their data available for research.

One of the first studies [17] on the NFT market, and the largest and most comprehensive to the date of publication of this thesis, gathered NFT transactions relative to purchases from multiple sources, namely CryptoKitties sales [119], Gods-Unchained [120], Decentraland [121], and OpenSea [122], together with data made available from Non-Fungible Corporation [123]. NFT transactions gathered this way present more attributes than Bitcoin

transactions. For instance, in addition to the timestamp, addresses involved and amount exchanged, they also refer to metadata about the NFT itself: the collection, item, encoded traits among others. The previously mentioned study by Nadini et al.(2021) exploited this wealth of data to extensively characterize the NFT market for the first time, showing the different categories of NFTs such as Art, Metaverse or Gaming, and how each contributed to the rise of the market. They showed how traders tend to specialize on similar objects and form cluster with other traders exchanging the same kind of NFTs. Exploiting the fact that NFT represent digital objects such as images, they also recovered the associated image files and were able to show how collections tend to contain visually similar NFTs, and how the visual features, together with the sale history, were the best predictors for the sale price of an NFT. In chapter 8 we will study a similar dataset, gathered from OpenSea, concerning 410 collections of NFT collectibles.

While the data about NFT transactions is generally richer than in Bitcoin, they still present several limitations which one has to remember when dealing with them. For instance, Ethereum shares with Bitcoin the same mechanism for addresses, meaning that a person can create a new address everytime they make a transaction. However, one can argue that with NFT transactions one is less motivated to do so, given the absence of illegality or police monitoring. Yet, someone can have several reasons why they might want to do so. For example, one might want to artificially increase the price of the NFT (price-pumping) [124], by repeatedly selling the NFT to themselves at higher prices, but using different Ethereum addresses. One other limitation to be aware of consists in only considering purchases, as transactions can also correspond to other actions such as minting the NFT or transferring to other addresses for other reasons than selling. Finally, gathering data from NFT marketplaces is increasingly harder, with platforms either not openly sharing data (e.g. SuperRare) or limiting access to it (e.g. OpenSea), making accessing data directly on the blockchain ledger the only, yet not so viable, alternative.

## Chapter 3

# Dark Web Marketplaces: reacting to the COVID-19 pandemic

COVID-19 gained global attention when China suddenly quarantined the city of Wuhan on January 23, 2020 [125]. Declared a pandemic by the World Health Organization on March 11, 2020, measures such as social distancing, quarantine, travel restrictions, testing, and tracking have proven vital to containing the COVID-19 pandemic [126]. Restrictions have shaken the global economy and reshaped the demand for goods and services worldwide [127]. Demand for many products has fallen; for example, the price of Brent crude oil decreased from 68.90 USD a barrel on January 1, 2020 to 43.52 USD as of August 2, 2020 [128, 129]. Meanwhile demand for other products, such as toilet paper [130], dramatically increased. As a result of increased demand, some products have been in short supply. Individual protective masks were sold in the United States at 7 USD on February 2, 2020 [131] and the price of alcohol disinfectant doubled on July 1, 2020 in Japan [132]. Additionally, anti-gouging regulations were introduced to control prices, which significantly affected the public attention on products related to COVID-19 [133]. As this trend has continued, further exacerbated by online misinformation, numerous customers have sought to fulfill their needs through illicit online channels [134, 135].

Dark Web Marketplaces (DWMs) offer access to the shadow economy via specialized browsers, like Tor [53], granting anonymity to its users also thanks to the use of cryptocurrencies for payments. DWMs offer a variety of goods including drugs, firearms, credit cards, and fake IDs [58]. Researchers have studied DWMs since the emergence of Silk Road [49], through a series of case studies [59, 60, 77], and comparative analyses [50, 51, 52, 78, 79, 80]. Past efforts have mostly focused on specific goods, such as drugs or digital products [56]. However, these studies experienced technical difficulties in data collection preventing researchers from analysing a large and up-to-date dataset. As a result, several questions remain open, among which are:

- how do DWMs react to sudden shocks (e.g., shortages) in the traditional economy?

- how do DWMs respond to trends in public attention?

In this chapter, we address these questions by analysing a new, large, and up-to-date dataset, comprising of 851,199 listings extracted from 30 DWMs between January 1, 2020 and November 16, 2020, right before the first worldwide vaccination campaign started in the United Kingdom [136]. In section 3.2.1 we studied the offer of COVID-19 related products on DWMs. We identified 788 COVID-19 specific listings that range from protective masks [137] to hydroxychloroquine medicine [138]. In Section 3.2.2 we investigated how this offer changed in time. We compared this COVID-19 related shadow economy with public attention measured through Twitter posts (tweets) [139] and Wikipedia page visits [140], finding correlation between the three time series, indicating how DWMs swiftly reacted to public attention. Finally, we inspected listings that mentioned delays in shipping or sales because of COVID-19, proxy for the indirect impact of COVID-19 on DWMs.

We significantly extended previous analyses that surveyed 222 COVID-19 specific listings extracted from 20 DWMs on a single day (April 3, 2020) [16] and offered a comprehensive overview of the DWM activity generated by the ongoing pandemic. We found that DWMs promptly respond to signals coming from the traditional economy, increasing or decreasing the offer of goods according to their availability on regulated markets. For example, protective masks appeared in DWMs in March, when they were in short supply in the regulated economy, and became more scarce on DWMs later on when masks could be easily bought in shops. We also found that DWMs swiftly react to changes in public attention as measured through Twitter posts and Wikipedia page views. Finally, we registered spikes in the number of listings mentioning COVID-19 in correspondence with lockdown measures in March and October. The work done in this chapter is based on publication [I].

## 3.1 Data and methods

### 3.1.1 Dark web marketplaces

The listings used for our study were obtained by web crawling DWMs. Web crawling consists of extracting data from websites and is performed by specialized software. Web crawling DWMs is a challenging task because crawlers must bypass several protective layers. Most DWMs require authentication and some even require a direct invitation from a current member. Additionally, strong CAPTCHAs [141] are implemented to avoid otherwise easy and automated access to the online DWM. Several research groups tried to overcome these challenges. Some examples are, HTTrack software used in [49], a combination of *PHP*, the *curl* library, and *MySQL* was proposed in [142], the Python library *scrapy* adopted in [143], and an automated methodology using the *AppleScript* language utilized in [144]. There are currently very few open source tools available [145, 141] for crawling DWMs, which often leaves companies and federal agencies to rely on commercial software [146]. Downloading content from DWMs remains a challenging task, and the objective becomes even harder when the research study requires monitoring multiple DWMs for an extended period of time.

Our dataset contains listings crawled from 30 DWMs between January 1, 2020 and November 16, 2020 by Flashpoint Intelligence [30], a company specializing in online risk intelligence. It includes the most popular DWMs in 2020, such as Hydra, White House, Empire and Dark Market [147, 16]. The crawling pipeline consists of authenticating into DWMs and downloading key attributes for each active listing. Each DWM was crawled for at least 90 different days. We categorized the COVID-19 specific listings into *PPE*, *medicines*, *guides on scamming*, *web domains*, *medical frauds*, *tests*, *fake medical records*, and *ventilators*. Representative examples of listings relative to these categories are presented in Table 3.1, with specific listing examples in Appendix A.2.1. Only a fraction of the selected listings were actual COVID-19 specific listings, since mitigation measures to prevent COVID-19 spreading have also impacted illegal trades of other listings. For instance, a vendor might sell cocaine and mention shipping delays due to COVID-19. We included such cases in the category COVID-19 *mentions*. For details about data pre-processing, see Appendix A.1, where we explain how we select listings related to COVID-19 and how we classify them in categories. We remark that our pre-processing pipeline is biased towards the English language, and this constitutes a limitation of our work.

Table 3.1: **Categories used to classify the selected COVID-19 dataset.** The first five categories constitute COVID-19 specific listings, while the last one, called COVID-19 *mentions*, includes listings mentioning one of the keywords in Table A.2 without selling actual COVID-19 specific listings.

Category	Examples
PPE	gloves, gowns, masks, n95
Medicines	azithromycin, chloroquine, azithromycin, favipiravir, remdesivir
Guides on scamming	how to illicitly get COVID-19 relief packages
Web Domains	covid-testing.in, coronavintheworld.com
Medical Frauds	antidotes, vaccines, allegedly curative recreational drug mixes
Tests	diagnosis, test
Fake Medical Records	medical record, medical certification
Ventilators	medical ventilators
COVID-19 mentions	computer, drugs, scam (excluding listings in the previous categories)

Overall, our dataset includes a total of 851,199 unique listings, which were observed a total of 8,538,593 times between January 1, 2020 and November 16, 2020. In Table 3.2 we report the breakdown of the number of unique listings and their total observations in each of the 30 DWMs. We did not find any mention of COVID-19 on 12 DWMs (Atshop, Black Market Guns, Cannabay, Darkseid, ElHerbolario, Exchange, Genesis, Mouse in Box, Rocketr, Selly, Skimmer Device and Venus Anonymous). This makes sense as these DWMs are primarily focused on specific goods with a pre-defined listing text structure. A brief description of each DWM together with its specialization can be found in Table A.4. On the remaining 18 DWMs, there were a total of 10,455 unique listings related to COVID-19, which constitutes less than 1% of the entire dataset. These listings were mostly composed of drugs that

reported discounts or delays in shipping due to COVID-19. Listings concerning more specific COVID-19 goods such as *masks*, *ventilators*, and *tests* were available on 13 DWMs (Connect, Cypher, DarkBay/DBay, DarkMarket, Empire, Hydra, MagBO, Monopoly, Plati.market, Torrez, CanadaHQ, White House, and Yellow Brick). There were 788 total COVID-19 specific listings in these DWMs which were observed 9,464 times during the analysed time period.

Table 3.2: **Information on data related to each DWM.** Number of days each marketplace was crawled, the number of unique listings, all and COVID-19 specific, and the number of listing observations, all and COVID-19 specific. CanadianHQ indicates The Canadian Headquarters marketplace.

Name marketplace	Days crawled	Listings All	Listings COVID-19 specific	Observations All	Observations COVID-19 specific
Black Market Guns	163	18	0	2,934	0
CanadaHQ	94	21,853	3	145,202	53
Cannabay	119	1,074	0	1,303	0
Cannazon	100	2,760	0	4,606	0
Connect	179	476	2	13,579	23
DarkBay/DBay	127	105,921	421	554,535	6570
DarkMarket	92	32,272	19	37,742	20
Darkseid	189	8	0	1,512	0
ElHerbolario	186	13	0	1,430	0
Empire	107	26,010	33	93,163	46
Genesis	188	216,792	0	2,174,217	0
Hydra	189	297	0	37,665	0
MEGA Darknet	135	754	0	1,596	0
Plati.Market	189	11,678	0	17,214	0
Rocketr	189	460	0	7,843	0
Selly	91	462	0	1,523	0
Shoppyy.gg	189	8,412	0	486,819	0
Skimmer Device	189	12	0	2,268	0
Tor Market	130	634	0	25,328	0
Venus Anonymous	177	84	0	14,644	0
White House	96	21,377	5	320,360	118
Willhaben	189	14,626	0	45,774	0
Yellow Brick	117	6,379	33	97,583	329
Total	> 90	472,372	518	4,088,840	7,159

### 3.1.2 Twitter

We sampled tweets related to COVID-19 using a freely available dataset introduced in Chen et al [139]. We downloaded the tweets ID from the public GitHub repository and then used

the provided script to recover the original tweets through the Python library *twarc*. We studied the temporal evolution of the number of tweets mentioning selected keywords, like “chloroquine”. In line with our dataset of DWM listings, most of the tweets considered were written in English and the time period considered ranges from January 21, 2020 to November 13, 2020.

### 3.1.3 Wikipedia

We used the publicly available Wikipedia API [140] to collect data about the number of visits at specific pages related with COVID-19, like chloroquine. The Wikipedia search engine was case-sensitive and we considered strings with the first letter uppercase, while the others lowercase. We looked for the number Wikipedia page visits in the English language from January 1, 2020 to November 16, 2020.

## 3.2 Results

We assessed the impact of COVID-19 on online illicit trade along three main criteria. First, we focused on the 13 DWMs containing at least one COVID-19 specific listing, analysing their offers in terms of the categories *PPE*, *medicines*, *guides on scamming*, *web domains*, *medical frauds*, *tests*, *fake medical records*, and *ventilators*, as introduced in Table 3.1. Second, we considered the 18 DWMs that included at least one listing in one of the categories in Table 3.1, thus adding listings to the COVID-19 *mentions* category in our analysis. We investigated the relationship between major COVID-19 events, public attention, and the time evolution of the number of active listings. Third, we quantified the indirect impact that COVID-19 had on all 30 DWMs under consideration by tracking the percentage of listings mentioning the themes of lockdown, delays, and sales. We linked their frequency to major COVID-19 events.

### 3.2.1 Categories of listings

Here, we focus on the 788 COVID-19 specific listings present in our dataset, observed 9,464 times in the considered time window. *PPE* is the most represented category, with 355 unique listings (45.1% of COVID-19 specific listings) observed 5,660 times (59.8% of observations of COVID-19 specific listings). The second most represented category is *medicines*, with 228 (28.9%) unique listings observed 1,917 (20.3% of all) times. Some *medicines* listings, which are often sold together, included 38 chloroquine listings, 65 hydroxychloroquine listings, 51 azythromycin listings and 45 Amoxicillin listings. Other *medicines* included 2 remdesivir listings, one of the drugs used to treat USA’s president Trump [148]. A breakdown of the *medicines* category together with a brief description of the specific drugs can be found in Table A.5, and multiple medicines are sometimes sold in the same listing. Another prominent category was *guides on scamming*, with 99 unique listings (12.6%). It includes manuals on how to earn money exploiting flaws in COVID-19 related government relief funds, and others

on how to exploit alleged pandemic related security weaknesses (e.g. online banking, delivery systems). A breakdown of the different kinds of guides can be found in table A.6. One DWM (MagBO) was specialised in the selling of web domains, like “coronavirusmasks.in,” with 50 unique listings (6.3%). Additionally, we classified 34 (4.3%) unique listings as *medical frauds*, which are listings that promised immunity from COVID-19 (no such product exists, at the moment of writing), or supposed devices able to detect COVID-19 in the air. These listings also included illicit drug mixes sold as cures. We also registered 17 tests (2.2% of COVID-19 specific listings), 3 *fake medical records* (0.4%) and 2 ICU *ventilator* (0.3%) listings. More details on these listings together with some examples are reported in Appendix A.2. There were a total of 252 vendors selling COVID-19 specific listings. Additionally, sellers posted multiple unique listings. In fact, 88 of them sold *PPE* (34.9%), 106 sold *medicines* (42.1%), 40 sold *guides on scamming* (15.9%), 15 *web domains* (6.0%), 23 sold *medical frauds* (9.1%), 13 sold *tests* (5.2%), 3 sold *fake medical records* (1.2%), and 2 sold *ventilators* (0.8%). The information in this paragraph is summarized in Table 3.3.

Table 3.3: **Summary statistics for the considered categories of listings.** For each category, we included the number of unique listings, observations, and vendors. If the same vendor sold listings in different categories, we counted it as one unique vendor.

Category	Unique listings	Total observations	Vendors
PPE	355 (45.1%)	5,660 (59.8%)	88 (34.9%)
Medicines	228 (28.9%)	1,917 (20.3%)	106 (42.1%)
Guides on scamming	99 (12.6%)	1,244 (13.1%)	40 (15.9%)
Web Domains	50 (6.3%)	189 (2.0%)	15 (6.0%)
Medical Frauds	34 (4.3%)	316 (3.3%)	23 (9.1%)
Tests	17 (2.2%)	51 (0.5%)	13 (5.2%)
Fake Medical Records	3 (0.4%)	9 (0.1%)	3 (1.2%)
Ventilators	2 (0.3%)	78 (0.8%)	2 (0.8%)
COVID-19	788 (100%)	9,464 (100%)	252 (100%)

It is important to note that vendors often do not provide complete information on their listings but rather invite direct communication to facilitate sales. In 391 (49.6%) unique listings, the vendor invited potential customers to communicate via email or messaging applications such as WhatsApp, Wickr Me, and Snapchat. Thus, 511 (64.8%) COVID-19 specific listings contained no information about the offered amount of goods, 579 (73.5%) did not provide shipping information, and 16 (2.0%) did not disclose the listing price.

*PPE* and *web domains* were the least expensive products with a median price of 5 USD. Followed by *medicines* with 33 USD, *guides on scamming* with 75 USD, *fake medical records* with 130 USD, *tests* with 250 USD, *medical frauds* with 275 USD, and *ventilators* with 1,400 USD. The distribution of prices for these categories can be found in Figure 3.1(a). It shows that many listings had a low price of around a few USD or less and only few listings exceeded thousands or more USD. The cumulative value of the detected unique listings was 563,202 USD, where we excluded listings with prices larger than 40,000 USD

using manual inspection. When vendors post listings at high price this typically indicates they have halted sales of an item with the expectation of selling it again in the future. We remove these anomalously high priced listings since they would largely overestimate the sales price of actually active listings [65]. The shipping information declared in the analysed listings involved a total of 18 countries or regions. Most of the vendors are willing to ship worldwide. Shipping from different continents appears possible because some vendors explicitly declare in listing descriptions that they have multiple warehouses across the globe, while shipping to any continent is done through specialized delivery services. The United States is the second largest exporter and shipping destination. The United Kingdom is the third largest exporter and importer, and no vendors explicitly mentioned Germany as a shipping destination even though it is the fourth largest exporter. Complete shipping information is available in Figure 3.1(b). Some examples of the COVID-19 specific listings are available in the Appendix A.2.1.




Figure 3.1: **COVID-19 related listings: price and shipping geography.** (a) Box plot of the distribution of listing prices for each COVID-19 category. The box ranges from the lower to the upper quartile, with the horizontal line indicating the median. The whiskers extend up to the 5<sup>th</sup> and 95<sup>th</sup> percentiles respectively. The dots represent outliers. (b) Shipping information in COVID-19 specific listings. Note that 545 (or 71.1%) of these listings did not provide any shipping information.

Figure 3.2(a) presents a word cloud built from the titles of the selected COVID-19 specific



the less represented DWMs are as shown in Table 3.2. The entire breakdown of the number of COVID-19 specific listings detected in each category is available in Figure 3.2(b).

In Figure 3.2(c), we ranked the DWMs by their share of vendors selling COVID-19 specific listings. The total number of vendors behind COVID-19 specific listings in our dataset is 252. Most vendors sold only one COVID-19 specific listing, while few of them sold more than ten different unique COVID-19 specific listings. In Appendix A.4, we analysed the distribution of COVID-19 specific listings for each vendor. We found that it was heterogeneous according to a power-law with an exponent equal to  $-2.3$  and 80% of the vendors had fewer than 5 unique listings, as shown in Figure A.1. This may imply that vendors of COVID-19 related products have a focus on a specific product category, or are just creating one-off listings to try to make quick money. In DarkBay/DBay, more than 15% of the vendors were selling COVID-19 specific listings, while in MagBO, The Canadian HeadQuarters, and Cypher this fraction was around 5% (with almost all other DWMs around 1%). This shows that law enforcement or intelligence intervention should not necessarily be approached evenly across marketplaces but instead focused on select marketplaces first with a higher concentration of COVID-19 specific listings. Finally, Figure 3.2(d) shows that essentially no vendors specialised on COVID-19 products, with only 7 vendors selling more COVID-19 specific listings than unrelated ones, 4 of which actually sold just one or two COVID-19 specific listings overall in our dataset.

### 3.2.2 Time evolution of DWM listings and public attention

The number of active unique listings evolved over time, as shown in Figure 3.3(a). The first COVID-19 specific listing in our dataset appeared on January 28, 2020, following the Wuhan lockdown [125]. In March, lockdowns in many countries [149, 150] corresponded to an increase in the number of these listings, whose number kept increasing until May. In June and July, when worldwide quarantine restrictions started to ease [151], we observed a decreasing trend in the selected COVID-19 specific listings, which continued until November. COVID-19 mentions followed the same trend with two notable exceptions. We observed two sudden increases in COVID-19 in correspondence of the second wave of contagions in Europe in September [152] and new lockdown measures in November [153]. Figure 3.3(b) shows the evolution of the total number of observed *PPE* and *medicines*, the two most available COVID-19 specific listings in our dataset (see Table 3.3). *PPE* followed a trend compatible with the overall observations shown in Figure 3.3(a), with a peak in May and a sudden decrease after July, as PPE have gradually become more available worldwide with respect to the shortage in the beginning of the pandemic. COVID-19 *medicines* remained approximately stable throughout these months, with a peak after USA president Donald Trump first referred to Chloroquine [154]. A different trend was found for COVID-19 *guides on scamming*, which saw spikes in the number of listings in correspondence to events related to relief program measures [155, 156, 157]. More details can be found in Appendix A.4, Figure A.2.

The time evolution of the listing prices followed a different pattern. We considered

the median price and its 95% confidence interval of active COVID-19 specific listings in Figure 3.3(c), and of active *PPE* and *medicines*, in Figure 3.3(d). Until March, the only COVID-19 specific listings concerned *medicines*, which influenced the overall median price. Afterwards, when *PPE* listings started to appear, they led the variation in the overall median price. In fact, over the entire time window, the median price of *medicines* listings was reasonably stable. *PPE* listings, instead, had a high price for March and most of April, possibly due to speculation. Interestingly, at the end of April, a vendor named “optimus,” active on DarkBay, started selling large quantities of *PPE* at 1 USD, putting many online listings at the same time, thus drastically reducing the median price, which remained low until July. Overall, “optimus” had 91 *PPE* listings during the registered period. *PPE* median price then increased back to the March level in July, when general worldwide availability of masks for the general population decreased the demand for small quantities of products. We report an analysis of the listings price for COVID-19 *guides on scamming* in Figure A.2 of Appendix A.4.




Figure 3.3: **Longitudinal analysis of DWM activity.** (a) Seven-days rolling average of active listings mentioning COVID-19 and COVID-19 specific listings. (b) Seven-days rolling average of the observed COVID-19 specific listings in the *medicines* and *PPE* categories. Black dashed vertical lines in panels (a) and (b) corresponded to significant COVID-19 world events, see Appendix A.3. (c) Seven-days median price with 95% confidence interval for COVID-19 specific listings. (d) Seven-days median price with 95% confidence interval for active COVID-19 specific listings in the *PPE* and *medicines* categories.

We also considered tweets and Wikipedia page visits as proxies for public attention, as already done in prior studies analysing the COVID-19 pandemic [158, 159, 160]. We compared trends in public attention with temporal variations in the number of active COVID-19 specific listings on DWMs. We focused our analysis on the *PPE* category and on relevant *medicines* in our dataset: hydroxychloroquine, chloroquine, and azithromycin. Figure 3.4(a) shows that a first peak in public attention on *PPE* was reached in late January following the Wuhan lockdown [125]. A second peak occurred in March [159] when *PPE* listings started to

appear in DWMs. The number of *PPE* listings reached their maximum in May. After May, *PPE* listings steadily decreased along with public attention. It is worth noting that May also marked the end of the first wave of contagion in many European countries [161]. *PPE* listings virtually disappeared in July, as products became more accessible in legal shops. On the contrary Twitter saw a huge spike in June, when many states decided to gradually lift lockdown measures [151], causing a public debates on mask wearing which increased the twitter signal to stable high levels until November.

A similar relationship between mass media news, public attention, and DWMs was registered for the listings regarding the three considered *medicines*, as shown in Figures 3.4(b) and (d). Four peaks in public attention were detected after four declarations from President Trump about these *medicines* [154, 162, 163, 164]. The number of active *medicines* listings closely followed. However, a closer look reveals the different shapes of the Wikipedia page visits, tweets, and DWMs curves. Wikipedia saw a very high peak of page visits after the first declaration from President Trump [154], and smaller peaks in correspondence in the following declarations. Tweets instead saw peaks of attention of increasing height. DWM listings on the contrary were much steadier in time and with little variation in the number of active listings throughout the first wave of the pandemic, while decreasing to a lower steady availability from the summer.




Figure 3.4: **DWMs and public attention.** (a)-(c) Seven-days rolling average of active listings selling *PPE*, together with the time evolution of the number of tweets referring to masks and of visits in the relative Wikipedia page visits. (b)-(d) Similar comparison as in panels (a)-(c) but considering active listings of hydroxychloroquine, chloroquine, and azithromycin. Black dashed vertical lines in panels (a) and (b) mark significant events related with COVID-19, see Appendix A.3. See Appendix A.4 for panels (a) and (b) with a linear y-axis.

### 3.2.3 Impact of COVID-19 on other listings

We considered the indirect impact of COVID-19 on all the 30 DWMs in our dataset. We analyzed all listings in these DWMs (COVID-19 related and beyond), and looked at listings mentioning: lockdown, using keywords “lockdown” or “quarantine,” delay, using “delay” or “shipping problem,” and sales, using “sale,” “discount,” or “special offer.” Examples of listings reporting these keywords are available in Appendix A.2.2.

Figure 3.5(a)-(b)-(c) shows the percentage of all listings mentioning these themes over time. The percentage of all listings in the 30 DWMs mentioning lockdown never exceeded 1%, as illustrated in Figure 3.5(a). It reached its maximum in November, when Europe started new lockdown measures [153]. Other peaks occurred in April and September, when nations first started to implement these measures [125, 149, 150] and at the beginning of the second wave of contagions in Europe [152], respectively. Delay mentions reached local peaks in March and May. These peaks occurred after major COVID-19 events, such as lockdowns [149, 150] and the situation in Europe starting to improve [161], respectively. Two global peaks, instead, were reached in September and November, when cases started to surge again in Europe [152] and when Europe started new lockdown measures [153], as shown in Figure 3.5(b). A similar pattern was visible for the percentage of all listings mentioning sales. In addition, we observed that sales had a first peak corresponding to the New Year, which is a common practice of many offline regulated shops, as displayed in Figure 3.5(c). Despite observing that the increase in the percentage of all listings mentioning sales, delays, and lockdown followed major events related to the pandemic, not all of these listings also mentioned COVID-19. We further researched this by plotting which percentage of the relative listings also mentioned COVID-19 in Figure 3.5(d). The percentage of listings mentioning that current sales were due to COVID-19 was less than 1%, while mentions of delays reached up to 40%. For lockdown it was 100%, as one can expect since lockdowns exist because of COVID-19. In the three selected cases, the percentages of listings mentioning COVID-19 followed the global awareness about the current pandemic: increasing trends from January to the July [125, 149, 150, 165], less attention during the summer [161], and a returning increase in September and November [152, 153].

## 3.3 Conclusions

We investigated the presence of listings related to COVID-19 in 30 DWMs, monitored over a ten-months period in 2020. We considered COVID-19 specific listings and COVID-19 *mentions*, found them in 13 and 18 DWMs, respectively. COVID-19 specific listings totaled 788 unique products and represented less than 1% of our dataset. The majority of COVID-19 specific listings offered *PPE* (45.1%), followed by *medicines* (28.9%), *guides on scamming* (12.6%), *web domains* (6.3%), *medical frauds* (4.3%), *tests* (2.2%), *fake medical records* (0.4%) and *ventilators* (0.3%). Most COVID-19 specific listings did not report the quantity sold (64.8%) or shipping information (73.5%). Almost half of these listings invited potential customers to communicate via email or messaging applications, like WhatsApp (49.6%).




Figure 3.5: **The indirect impact of the COVID-19 pandemic.** Percentage of all active listings mentioning the themes lockdown, delay and sales in panels (a), (b), (c), respectively. (d) Percentage of active listings in panels (a), (b), (c) that mentioned also COVID-19 in their listings. Black dashed vertical lines in panels (a), (b), and (c) corresponded to major COVID-19 events, see Appendix A.3.

Although direct communication fosters a trustworthy vendor-buyer relationship and may lay the ground for future transactions outside DWMs, it also exposes users to higher risk of being traced by law enforcement [166].

In our dataset, DarkBay/DBay is featured prominently among DWMs offering COVID-19 specific listings. Ranking in the top 100 sites in the entire dark web [167], DarkBay/DBay is regarded as the eBay of the dark web because it offers more listings categories than other DWMs [168]. It was also frequently accessible during the period of time monitored during this research, with an uptime of 80%, higher from the 77% uptime of Empire, the largest global DWM at the time of writing [169].

Our work corroborates previous findings and expands them in several ways. To the best of our knowledge, the most extensive report to date examined the presence of COVID-19 specific listings in 20 DWMs on one single day (April 3, 2020) [16]. Despite only a subset of overlapping DWMs between that report and our study, (Cypher, DarkBay/DBay, DarkMarket, Empire, Monopoly, Venus Anonymous, White House, and Yellow Brick) we both assessed that COVID-19 specific listings constituted less than 1% of the total listings in the DWMs ecosystem. These listings were mostly *PPE*, followed by *medicines* and they were found in only a few DWMs, while non COVID-19 specific listings were widespread.

An important novelty of the present study is the analysis of the temporal evolution of DWM behaviour and its relationship to public attention, as quantified through tweets and Wikipedia page visits. Following the Wuhan lockdown [125], we observed a first peak in

public attention [160], and a corresponding emergence of the COVID-19 specific listings. A second peak in public attention occurred in March, when quarantine measures were adopted by many European countries [149, 150]. Again, during the same period, the number of COVID-19 specific listings sharply increased. When worldwide quarantine began to ease [151] in many countries, in June and July, we registered a decrease in public attention and in available COVID-19 specific listings. Towards the end of 2020, we did not detect significant variations in COVID-19 specific listings and public attention, in correspondence of the second wave of contagions [152] and new lockdown measures in Europe [153]. Both vendors of COVID-19 specific listings and public attention have adapted to the COVID-19 pandemic and react more smoothly to its development.

Listing prices correlated with both variations in public attention and individual choices of a few vendors. Median price experienced a sharp increase in March, probably due to speculation, and then decreased in April due to the choice of a single vendor responsible for 91 listings, named “optimus.” The vendor sold a large quantity of *PPE* at 1 USD only, which constituted the 37% of active *PPE* listings in April. Finally, we observed an increase in the percentage of all listings citing delays in shipping and sale offers, which peaked in March, May, September, and November. Similar to a prior work that found Wikipedia page visits of a given drug to be a good predictor for its demand in DWMs [170], we provide further evidence that the DWMs ecosystem is embedded in our society and responds in line with social changes [171]. The DWMs ecosystem swiftly reacted to the pandemic by offering goods in high demand, and even offering vaccines already in March, when no tested vaccination existed.

Our research shares some limitations with previous studies, namely that not all active DWMs were surveyed. For instance, we did not analyse 12 of the DWMs explored in the previous report [16]. It must be noted, however, that the number of active DWMs is constantly changing due to closures or new openings [8] and obtaining full coverage is challenging due to the active efforts of DMWs to obstruct research studies and law enforcement investigations, for example through the use of CAPTCHAs. Another limitation is the lack of reliable fully automated annotation method: this forced us to manually annotate listings and thus limited our analysis to listings only directly related to COVID-19. One key problem to be solved in this regard is the presence of false positives when doing a keyword search.

## Chapter 4

# Dark Web Marketplaces: adapting to the COVID-19 pandemic

COVID-19 has caused a worldwide economic and public health crisis, that demanded and stimulated a global response. Hundreds of possible COVID-19 vaccines have been proposed [172] since the first officially approved vaccines in late 2020, like Sputnik [173] and Pfizer/BioNTech [174, 175, 176]. The subsequent initial scarcity and unequal distribution of COVID-19 vaccines [177] have generated concerns about illicit trade early on. Interpol warned about illicit offering of COVID-19 vaccines already on December 2, 2020 [178], while Europol confirmed the sale of fake COVID-19 vaccines on dark web marketplaces (DWMs) on December 4, 2020 [179], warning that it “may pose a significant risk to public health”. Understanding how DWMs reacted to the demand for vaccines is therefore crucial to allow policy and public health agencies to be prepared and effectively counteract these threats in the future.

Interpol and Europol’s concerns were validated by early research showing that DWMs have been an important channel to access online illicit trade during the pandemic, with masks, COVID-19 tests, and alleged medicines consistently advertised on these platforms. In a first report [16], 222 COVID-19 related unique listings were registered on April 3<sup>rd</sup>, 2020 in 20 DWMs. In the previous chapter we also showed how 788 COVID-19 related listings were observed 9,464 times between January 1, 2020 and November 16, 2020 in 30 DWMs, showing how DWMs swiftly reacted to shortages and public attention by offering sought-after products like masks and hydroxychloroquine. More recent reports, carried by the Global Initiative and Europol, have suggested that the overall structure of illicit online trading has gained significant benefits from COVID-19 [134, 135].

Here, we report on our analysis of 194 DWMs until July 22, 2021. In doing so we extend the analysis of the previous chapter, focused on the immediate reaction of DWMs to the shock caused by the onset of the COVID-19 emergency [14], to consider how DMWs have responded and further adapted to the ongoing pandemic. Furthermore, the period we cover

includes the milestones of COVID-19 vaccines being approved and made available, allowing us to investigate their offer on unregulated markets. We detected a total of 10,330 unique listings that were directly affected by COVID-19, i.e., mentioning COVID-19 either in their body or title. In section 4.2.1, we show how, among these listings, 248 were offering vaccines. It is important to note that a listing does not correspond to the sale of a unit, as sometimes happens for example on Ebay, but corresponds to the availability of multiple units of a product, similarly to what happens for example on Amazon. Listings related to approved vaccines were initially detected on the Invictus marketplace starting from November 17, 2020, almost 2 weeks before their official approval. Also, listings offering a fabricated proof of vaccination were registered on the Hydra marketplace since February 15, 2021. In section 4.2.2 we show how these listings replaced previously identified COVID-19 related products, like PPEs, COVID-19 tests, and guides on how to illicitly obtain COVID-19 relief funds. The availability of these products have decreased with respect to previous observations, with only 187 listings detected between November 2020 and July 2021 against the 788 registered between January and November 2020 [14]. Many vendors selling these products are highly specialised in only a type of product and willing to ship worldwide, thereby increasing the number of potential customers. Finally, by analysing all listings mentioning COVID-19, in section 4.2.3 we assess the overall impact of COVID-19 on DWMs. We show that drugs are the only traditional DWMs product to have been indirectly, and increasingly, affected by the pandemic, with vendors mentioning both pandemic related supply issues and delays. This chapter is based on publication[II].

## 4.1 Data and methods

Our dataset includes the most popular DWMs in 2020 and 2021, such as White House, Empire, Hydra, and DarkMarket [147, 16] and was gathered by Flashpoint [30]. It consists of an extension of the dataset used in chapter 3, covering more markets and a longer period of time. Note that the landscape of active DWMs is constantly changing: Empire exit scammed, meaning that it closed down without any notice and taking away the deposited funds of its users, on August 23, 2020 [180], while DarkMarket was shut down by Europol on January 12, 2021 [181].

Our DWMs dataset is used to complement and extend the analysis we previously performed for the period between January 2020 and November 2020 [14]. The new dataset also covers the period following the approval of the vaccines and their actual distribution to the population, i.e. Nov 2020 to July 2021, and allows us to observe the evolution of COVID-19 related products over the second part of the pandemic. We also add several new DMWs, increasing their number from 30 to 194, and comprehending a total of 10.8 million unique listing titles. Only 84 of these DWMs mentioned COVID-19, 20 DWMs offered COVID-19 related products, and 19 vaccines, see Table A.4. Each unique listing is observed at most once per day.

During the considered period of the COVID-19 pandemic, the illicit offer of vaccines

(i) Vaccine related set of keywords
antibod, vaccin, antidot, vacun, immun, Инокул, вакцин, прививк, Ревакцин, Инокул, 疫苗, 反, impfstoff, Gegenmittel
(ii) COVID-19 and brands related set of keywords
covid, corona, ковид, Коронавирус, Пандеми, Вирус, Спутник V, Инфекци, Симптом, 新冠病毒, 武汉肺炎, couronne, pfizer, astrazeneca, moderna

Table 4.1: **Search of COVID-19 vaccines.** Keywords used to pre-select vaccine listings from the original dataset. Words are truncated to include different suffixes (e.g., vaccin yields vaccine, vaccination, vaccinate, etc.)

Category	Keywords
Guides on scamming	guide, fraud, exploit, scam, loan, relief, scampage, cashout
Medicines	chloroquin, azithromycin, favipiravir, ritonavir, lopinavir, remdesivir, dexamethasone, ciprofloxacin, doxyciclin, oseltamivir, metronidazol, ivermectin
PPE	mask, glove, gown, surgical, sanitiser, sanitizer, ppe
Test	test kit, covid test, pcr test, antigen test, corona test, diagnostic, diagnosis
Web domain	https, www., http://, .com, .co.uk, .dk, .org, .info, .in, .net

Table 4.2: **COVID-19 products related keywords.** Keywords used to pre-select listings selling COVID-19 related products before their manual annotation, organised by category.

constituted one of the biggest threats for global public health. We therefore use a method to detect vaccine listings that ensures the highest possible coverage and accuracy. From the listings, we considered two different text fields: the title and the body (that is, the listing’s detailed description). We then pre-selected all listings which contained, either in the body or title, at least one word from two different lists of keyword. These lists of keywords are shown in Table 4.1: the first list contains keywords related to vaccines; the second list contains keywords related to COVID-19 or vaccine brands like Pfizer/BioNTech. Note that using keywords like “antibod” or “vaccin” allows us to match all words including these sets of strings, such as, antibode, antibodes, vaccine, vaccines, vaccinations, and so on. We considered several different languages, such as, English, Russian, Chinese, and German. Afterwards, we manually inspected the listings to exclude false positives from the dataset; we categorized the listings in specific subcategories (e.g. specific brands), and we standardised the analysed attributes for the analysis. For example, we converted all prices to USD at the daily exchange rate at the time of observation.

Such a method is not feasible as more products are searched, because the number of listings to be manually annotated is too large. As already done in our previous work [14], we then limit our analysis to all listings mentioning COVID-19, using one of the follow-

ing keywords: “corona virus”, “covid”, “coronavirus” either in the title or description. To analyse COVID-19 related products, we first pre-selected a subset of these listings mentioning keywords in specific categories, see Table 4.2, and then manually annotated these listings. With respect to our previous effort, we find a new product category, which we call *malware*, while no listing in the *ventilator* category was found. Then, we characterize all listings mentioning COVID-19 by means of Natural Language Processing techniques. We perform such analysis on the title, which contains essential information about the listing. First, we use doc2vec, a deep learning model that creates vector embeddings of sentences and paragraphs, to map the listing titles into high-dimensional numerical vectors. In particular, we make use of the specific “*paraphrase-mpnet-base-v2*” model implemented by the python package *sentence-transformers* [22], a pre-trained model which embeds sentences into a 768-dimensional vector space. In this space, semantically similar sentences are mapped in vectors close to each other, allowing for a quantitative way of detecting similar listings. In order to capture clusters of similar titles, we first need to reduce the dimension of the space without losing the information encoded in the distance between vectors. To this end, we use the UMAP algorithm [23] to map the 768-dimensional vectors to a 2D space, preserving its structure. We finally employ the *hDBSCAN* algorithm [182] (with minimum cluster size of 100 documents) to cluster these 2D vectors. We then label each cluster according to the category of products the listings refer to, manually inspecting the highest ranking words as ranked by the *tf-idf* algorithm, a statistical measure that evaluates how relevant a word is in a collection of documents. When not otherwise specified, we used default parameters.

While each listing had an associated url to determine its uniqueness, which allowed us to track listing over time, vendors receiving bad reviews sometimes put identical copies of the same listing online. To overcome this issue and correctly count the number of listings, we created a new identifier of unique listings. We considered two listings as unique if the same vendor was posting listings in the same market, having only small variations in the title. We also excluded listings with prices larger than 40,000 USD. Vendors post listings at high price to hold sales of these relative items, with the expectation of offering it again in the future [65].

## 4.2 Results

### 4.2.1 Vaccine listings

We start by analysing COVID-19 listings since November 2020. We found 248 unique listings offering vaccines and manually categorised them in three categories: *approved vaccines*, *unspecified vaccines* and *proofs of vaccination*. Listings in the *approved vaccines* category explicitly mentioned official vaccines, an example being the Pfizer/BioNTech vaccine that was offered at 500 USD on the Invictus marketplace, see Fig B.1. Listings in the *unspecified vaccines* category instead referred to unbranded vaccines, for example by offering alleged unapproved vaccines well before official clinical trials were completed, as shown in Table B.1.

For instance, our previous analysis [14] found 34 listing advertising fake cures for COVID-19, including antidotes, vaccines, and allegedly curative recreational drug mixes. These listings were scam, since no official vaccine was approved in the considered time period. Listings in the *proofs of vaccination* category offered a fabricated certificate of COVID-19 vaccination, as the fake COVID-19 passport offered at 55 USD on the Hydra marketplace, see Fig B.2 with its English translation in Table B.2. The *unspecified vaccines* category contained 94 listings, followed by the *proofs of vaccination* category with 80 and then the *approved vaccines* category with 74 listings. The *unspecified vaccines* category also has the highest number of vendors, with 61 offering these products across 13 different DWMs. Similar statistics for the other categories can be found in Table B.7.




Figure 4.1: **DWMs and COVID-19 vaccines.** (a) Number of unique listings offered in each DWM. “BB. house” stands for Big brother house, while “Canadian HQ.” to The Canadian Headquarters. (b) Violin plot of the prices in USD at which vaccines were offered, showing the distribution of prices for the three categories. (c) Number of listings indicating where vaccines are declared to be shipped from and to. “N.A.” stands for not applicable and “Russia” for Russia and Eastern neighbouring countries. (d) Number of vendors offering a vaccine in a given category. Only the lower triangle of the matrix is shown because it is symmetric, where its diagonal represents vendors offering only listings in that category.

In Fig. 4.1 we characterize the offer of these listings. We start by considering how the offer of vaccines was distributed across markets. The majority of vaccines were offered in the Agartha marketplace, with 108 listings, followed by Hydra with 67, which offered 65 out of the 80 fabricated COVID-19 vaccination certificates in our dataset. Fig 4.1(a) shows the category of listings offered by each DWM with at least one vaccine. 11 of these DWMs are specialized in offering only one category of listings, with one DWM only offering *approved vaccines*, 5 DWMs only offering *unspecified vaccines*, and 5 DWMs *proofs of vaccination*.

Three DWMs, Agarth, Liberty, and Yakuza, offer at least one listing in each of the three categories considered. The DWMs specialization can be seen in Fig B.5. Vaccine listings have a short lifetime on a DWM, with most listings that are offered for less than 25 consecutive days, see Fig B.4. Such short lifetimes may be due to platform moderation, which in some cases explicitly prohibit such listings, supplies running out or even vendors taking down the listings because of bad reviews. However, such claims are not verifiable with our current dataset.

Regarding the price of vaccine listings, Fig 4.1(b) shows its distribution in the three categories under consideration. Listings in the *approved vaccines* category have prices ranging from 40 to 2,400 USD; listings in the *unspecified vaccines* category between 25 USD to 6,060 USD; and listings in *proofs of vaccination* category from less than 1 USD up to 814 USD. Proofs of vaccination were the cheapest products, probably because they consist of fake documentation (e.g., falsified COVID-19 passport). Price of *approved vaccines* listings varied depending on the vaccine brand offered, see Fig B.7. The first listing in this category to be offered was the Pfizer/BioNTech vaccine at 1,000 USD. The other 44 listings offering the Pfizer/BioNTech vaccine proposed prices ranging from 200 to 2,400 USD. The Astrazeneca/Oxford vaccine, the second to be officially approved, was offered on DWMs since December 27, 2020. Only four listings offered this vaccine, ranging from 300 to 900 USD. The other approved vaccines offered on DWM were Moderna with 21 listings, Johnson&Johnson with four, Sputnik V with four, and Sinopharm with two. Their prices ranged from 40 to 2,000 USD. We speculate that one possible reason behind the skewness of these price distributions could be the presence of scam listings pretending to be selling these products at very low prices.

A natural next step is to analyse the geography of this trade, which we can do by looking at the shipping origin/destination information advertised in the listings. Most vendors declared that they would ship anywhere in the world, a behaviour that facilitates illicit trade. Vaccine warehouses were mostly in USA, followed by Germany and UK. Also, many listings do not declare any shipping information and all general shipping statistics are visible in Fig 4.1(c). In the 58% of the cases where no shipping information is declared, vendors invite potential customers to a direct interaction through Whatsapp, email, or phone. The percentage of listings where vendors suggest to initiate a direct interaction varies depending on the category considered. It happens for 78 (or the 84%) of listings in the *unspecified vaccine* category, for 64 (or the 85%) of listings in the *approved vaccine* category, and for only three listing in the *proofs of vaccination* category. This last low number is due to Hydra marketplace, which sells 64 proofs of vaccinations but whose vendors never shared their contact information.

Do these vendors sell multiple kinds of products related to vaccines? Or do they focus on a single category? Fig 4.1(d) shows that vendors offering *proofs of vaccination* were specialised, with only one vendor also offering *approved vaccines* and two unspecified vaccines. On the contrary, 11 vendors were offering both *vaccines* and *unspecified vaccines*. We did not observe any vendor offering listings in all three categories. Moreover, most vendors (tracked

by username in the absence of PGP signatures) offer only one COVID-19 listing and trade in only one DWM, with the notable exception of a vendor, who had twelve listings in eleven different DWMs, as detailed in Fig B.6.




Figure 4.2: **Temporal evolution of COVID-19 vaccine listings.** (a) Cumulative number of listings over time in the three categories considered. (b) Average price over time in the same three categories, computed with a 90-days moving window. Vertical dashed black lines represent relevant pandemic events.

We now consider the time evolution of this offer of vaccine listings on DWMs, as shown in Fig. 4.2. The evolution of vaccines on DWMs closely followed major COVID-19 related events, as shown in Fig B.3 and a sample of which is also shown on the background of Fig 4.2(a). Fig 4.2(a) shows that multiple vaccine listings were simultaneously present on DWMs when the first vaccination trials were undergoing, between March 16, 2020 and April 14, 2020 [183]. No more vaccine listings were present on DWMs from July 1, 2020, coincidentally with the end of the first wave of contagions in Europe (June, 2020 [151]). These listings reappeared on September 16, 2020, at the beginning of the second wave of infections that started in September 2020 [152]. Up to that moment, we detected only COVID-19 listings in the *unspecified vaccines* category. The first listing in the *approved vaccines* category was a Pfizer/BioNTech vaccine and was offered since November 17, 2020, two weeks before its first official approval on December 2, 2020 by the UK regulator MHRA [174]. A similar pattern was registered for the first AstraZeneca/Oxford vaccine listing on DWMs. It was offered on December 27, 2020, three days before the first official approval of this vaccine (by the UK) on December 30, 2020 [184]. The remaining approved vaccines, Johnson&Johnson, Moderna, Sputnik V, and Sinopharm, all appeared in the first half of March, when we started to monitor the Agarthra marketplace. All *approved vaccines* listings disappeared on DWMs after May 1, 2021, albeit there may be other DWMs offering these products that are not part of our analysis. Since listings in the *unspecified vaccines* category continued to be observed until July 2021, we speculate that vendors were starting to have multiple vaccine brands, and they did not specify anymore which one they are selling. For more details, see Fig B.8(a). Listings in the *proofs of vaccination* category emerged on February 15, 2021, when airlines were encouraging governments to allow certificates of vaccinations to become a way to safely travel [185].

Finally, we looked at the temporal evolution of the average price of these listings. The three categories followed different trends, as visible in Fig 4.2(b). The price of *unspecified vaccines* was high between March and May 2020, when DWMs vendors likely tried to profit from the initial lack of COVID-19 medications [14]. Afterwards, their mean price has gradually decreased, meaning that the new listings appearing on DWMs were offered at progressively lower prices. However, the average price rose back to March levels in January 2021, when vaccinations campaigns around the world were starting. The availability of officially tested vaccines led to the emergence of listings advertising officially approved vaccines on DWMs since November 2020. The average price of these listings have floated over time between a few hundreds USD to more than a thousand. For more details, see Fig B.8(b). Finally, the needs for a certificate of COVID-19 vaccination had meanwhile increased, and so had the price of listings in the *proofs of vaccination* category. Vaccines certificates have gradually become mandatory in many countries, and especially for international travel, and their sale on DWMs confirms what researchers had hypothesised [186], warning against similar situations happening in the future.

#### 4.2.2 Other COVID-19 related products

DWMs have been a venue for the sale of other licit and illicit COVID-19 related products, like PPEs, tests, or medicines, as reported for the period from January to November 2020 [14]. Here, we monitor COVID-19 related products in the second part of the pandemic, between November 2020 and July 2021, see Table B.8 and Fig 4.3. Listings are divided in six different categories: *PPEs* represent healthcare objects like masks; *medicines* COVID-19 related medicines like hydroxychloroquine; *guides on scamming* are instructions on how to get relief funds; *tests* represent COVID-19 tests; *web domains* that are related to COVID-19 like "covidtest4you.com"; and *malware* represents malicious software to hack COVID-19 test or vaccination records software. Listings from these categories are offered in 21 DWMs, and are available in multiple markets. *Malware* and *web domains* are an exception because they are sold in two specific markets only. We find that *PPEs* and *medicines* have almost disappeared from DWMs w.r.t. previous observations [14, 16]. *PPEs* listings are mostly advertising bulk sales at high prices, coherently with the end of shortages of these products, while *medicines* listings, like hydroxychloroquine, are substituted by vaccines and present on DWMs in a lower number, with only 3 listings advertising Ivermectin [187]. On the contrary, *guides on scamming* were still present with comparable numbers, claiming to teach ways to access COVID-19 relief funds in different countries. Notably, the number of listings offering COVID-19 tests had also increased, with tests increasingly being required for travel or work. We also found 4 listings advertising malware to illicitly access official systems to record test results or even vaccinations.

Fig 4.3(a) shows the distribution of unique listings on each DWM offering them. These listings are concentrated in 4 DWMs, with the majority of them offering less than 5 listings. However, there is less category specialization w.r.t. what we observed for the vaccines, with multiple markets offering listings in different categories. Prices are also very heterogeneous,




Figure 4.3: **COVID-19 related products.** (a) Break-down of COVID-19 related products by category and market. (b) Boxplot representing the price distribution of listings in each category. Horizontal lines represent the median value, box ends the first and third quartiles, and whiskers minimum and maximum values, respectively. (c) Number of listings indicating where COVID-19 related products are declared to be shipped from and to. “N.A.” stands for not applicable and “Russia” for Russia and Eastern neighbouring countries. (d) Number of vendors offering a COVID-19 related product in a given category. Only the lower triangle of the matrix is shown because it is symmetric, where its diagonal represents vendors offering only listings in that category.

with *test*’s median price highest at a few hundreds USD, and *web domain*’s lowest at just 4 USD, but also listings inside the same category ranging from 10 USD to 1000 USD in all categories but *web domain*, see Fig 4.3(b). In Fig 4.3(c), we show the origin and destination of the considered listings, as declared by vendors. The majority of listings declare to be shipping worldwide, while the United States is the country appearing the most as declared origin of the listings. Russia and Eastern neighbouring countries are both origin and destination, mainly because of *proof of vaccination* listings offered on Hydra, whereas UK, Germany and India appear almost only as countries of origin. Other countries/regions are declared, but less frequently. Fig 4.3(d) shows vendors specialization regarding COVID-19 related products. All categories show highly specialised vendors, except *PPE*, where only one vendor out of seven sells only in that category, and *tests*, where less than 55% of vendors sell only such products.

### 4.2.3 Listings with COVID-19 mentions

In this section we extend our analysis to the offerings of products that mention Covid-19 in the title or body, without being directly related to pandemic-products, thus providing a richer assesment of the overall impact of COVID-19 on DWMs. We extend our previous analysis by considering listings appearing until July 2021 and by categorising the selected listings, providing a richer and deeper picture of how DWMs were indirectly affected by COVID-19.

We characterise products mentioning COVID-19 with state-of-the-art Deep Learning based Natural Language Processing techniques [188, 23, 182], see Methods for more details. As shown in Fig 4.4(a), we find 13 different categories of listings corresponding to different kinds of products. In addition to the already discussed COVID-19 related products, only drugs appear to be mentioning COVID-19, while other traditional DWMs’ products like stolen IDs or credit card dumps don’t, showing which kind of goods reacted to, or where affected by, the pandemic. We then analyse the temporal evolution of these categories. We show the number of active listings for 4 large categories in Fig 4.4(b), while all other categories are shown in Fig B.9. Drugs show an overall increasing trend throughout the whole period. Different categories, however, display different fluctuations in time, showing how different goods behave in an heterogeneous way with respect to COVID-19. For example, at the end of our covered period we can see *thc* and *psychedelichs* showing a flat trend, while *cocaine* and *mushrooms* are increasing.

While it is not possible to understand the reasons behind each single temporal trend, we can gain more insights on why drugs are increasingly mentioning COVID-19 by investigating which themes are recurrent in these listings. In Fig 4.4(c), we count mentions over time of three different sets of keywords, which can be used as general proxies of the indirect impact of the pandemic. We considered (i) *lockdown*, by tracking listings mentioning “lockdown” or “quarantine,”; (ii) *delay*, by monitoring listings using keywords “delay” or “shipping problem,”; and (iii) *sales*, by searching for “sale,” “discount,” or “special offer.” For instance, sellers may mention lockdown to justify the lack of international deliveries or problems with their supply. Similarly, sellers may mention possible delays due to COVID-19 related restrictions and supply issues, or promote sale during the economically-challenging COVID-19 period [14, 189]. *Lockdown* mentions are always lower than the other two themes, peaking in summer 2020 but staying always lower than 20%. *Delay* mentions instead rapidly increase during the first months of the pandemic, and have been oscillating around 60% of the listings since then, showing how drug vendors have been warning about possible delays throughout the whole observed period, confirming what’s already been independently shown for the first phase of the pandemic [189]. Finally, *sale* mentions show larger fluctuations between as low as 15% to even as high as 80%. In particular, we can observe peaks related to the pandemic at the beginning of key COVID-19 related events: lockdowns in March/April 2020, in Summer 2020 coincidentally with openings in the western world, in October 2020 when the second wave started hitting Europe, and in February 2021 when the Delta variant started spreading in the world. By looking at mentions of these themes across all listings in our dataset, we find

that overall mentions of *lockdown*, *sales*, and *delay*, have decreased since the beginning of the pandemic, validating our finding that drugs-related listings are the product most impacted by COVID-19, see Fig B.10.

Automatic keyword search has allowed us to uncover macroscopic trends, but it fails to capture finer details which can only be uncovered by in-depth looks at the texts of the listings. We therefore resort to a qualitative analysis of their descriptions. First, we already noticed that mention of delays in drug listings are still frequent, amounting to 56 % of the listings. While vendors generally preemptively mention possible delays due to COVID-19, we find numerous mentions of USA based vendors blaming USPS for these, as shown in one example reported in Table B.3: "THE USPS IS UNDERFUNDED AND MAY BECOME UNRELIABLE COMPARED TO THE PAST! (ESPECIALLY DURING COVID-19 AND HOLIDAYS!)". These claims reflect widely reported issues with the United States Postal Services since June 2020 [190]. Moving away from delays, we find that 10% of vendors mention COVID-19 by ensuring potential clients that they are taking all necessary safety measures when preparing the deliveries. An example of this is reported in Table B.4. Finally, we find listings mentioning limited stocks due to the pandemic, as shown in Table B.5, where the vendor claims that "stocks are almost exhausted by Corona Covid 19".

### 4.3 Discussions and conclusion

In this chapter we have studied if, and how, DWMs have adapted to the ongoing COVID-19 pandemic along multiple dimensions, with a special focus on the sale of COVID-19 vaccines. Covid-19 vaccines have indeed been a key element of the exit from the emergency phase of the pandemic, and regulators and international agencies warned of their possible illicit trade on DWMs and the associated health risks [178, 179]. The covered period, ending at the end of July 2021, included the second phase of the pandemic, i.e., when vaccines became available. We have identified a sharp increase in the number of listings selling vaccines and proofs of vaccinations, from 34 between January and November 2020 (when no vaccine had been released yet) to 248 after, including officially approved vaccines like Pfizer or Moderna. Vaccine related listings have replaced other previously observed COVID-19 related products (e.g. PPE and hydroxychloroquine), whose presence has been steadily decreasing since November 2020. While assessing the overall COVID-19 impact through the analysis of listings explicitly mentioning COVID-19, we have found that drugs were the most affected traditional DWMs product. Our results extended previous analyses [14, 16] on the impact of COVID-19 on DWMs both in terms of duration of the monitored period and of breadth of the analysed products.

A key contribution of the present work is the study of the interplay between DWMs and the COVID-19 pandemic, after the official approval of vaccines and when the pandemic had been present for a long period of time. It was shown in chapter 3 that, at the onset of the pandemic, when a product was in shortage in the regular economy, or public attention was focused on it, listings advertising its sale appeared on DWMs. For instance, this is what hap-




Figure 4.4: **Characterization of COVID-19 mentions.** (a) UMAP representation of doc2vec embeddings. hDBSCAN clustering finds 13 meaningful categories covering COVID-19 related products and all major drugs sold on DWMs. (b) 30-day rolling average of active listings in 4 categories of listings mentioning COVID-19: *thc*, *psychedelics*, *cocaine*, and *mushrooms*. (c) 30-day rolling average of fraction of previously identified drugs listings mentioning 3 different COVID-19 related themes: lockdowns, shipping delays and sales, with key pandemic events highlighted with vertical dashed black lines.

opened for PPE and hydroxychloroquine during the first phase of the pandemic. Since in the currently observed period of time these products were easily available on regulated markets, we coherently detect that these products disappeared in the second phase of the pandemic. In late 2020, we have seen the same pattern with vaccines, which appeared around the time of their official approval, reflecting the claims of other mass media news [191, 192, 193]. They then spiked at the beginning of 2021, to be later replaced by fabricated proofs of vaccinations with the increasing need of vaccine passports and green passes [194, 195]. Mentions of lockdowns, delays, and sales followed similar dynamics, with spikes observed in the first phase of the pandemic in 2020 and their mentions steadily decreasing during the second phase. However, we found that drugs listings mentioning COVID-19 increased in time, with numerous mentions of delays and sales, some of which are related to stock shortages and increase in health security measures, as unveiled by our qualitative analysis. Our results confirm what was already observed for other external shocks creating extraordinary demand for specific goods. For instance, it has been shown that the restriction of access to hydrocodone combination products, the most commonly prescribed opioid, in the United States in 2014

caused DWMs to step in to meet the unaddressed demand [196, 197].

A limitation of the present work is that, while the number of DWMs simultaneously monitored over time is greater than most previous studies, we cannot ensure that all DWMs were surveyed. In fact, the number of active DWMs is constantly changing due to closures or new openings [8] and obtaining full coverage is challenging due to the active efforts of DWMs to obstruct research studies and law enforcement investigations. Moreover, our study is limited to what vendors advertise on these platforms, as we have no data on actual purchases to quantify how many people have been endangered by this phenomenon. Future work, relying on backend servers seized during police takedowns of DWMs, could improve on our study by overcoming these limitations.

The diffusion of illicit vaccines on DWMs, together with the simultaneous decrease of PPEs and medicines, confirms the link between product shortages, public attention and listings on DWMs. This phenomenon has the potential to pose a serious public health threat, as DWMs have become increasingly easier to access, resilient to police closures [8] and shown to be a catalyst for decentralized peer to peer trading between buyers and sellers of illicit items [198]. The purchase of unregulated, and possibly fraudulent, health related items on DWMs poses a concrete health risk for the buyers. Our results call for more investigation of DWMs to anticipate such dangers in future public health crisis.

## Chapter 5

# Macroscopic properties of buyer-seller networks in online marketplaces

Much of online trade happens on regulated and unregulated online marketplaces. Regulated online marketplaces include Amazon, Craigslist, eBay, Walmart, Alibaba (China), Rakuten (Japan), Gumtree (UK), and Mercado Libre (South America). Unregulated online marketplaces, such as Silk Road, AlphaBay, and Hydra, that specialise in the sale of illicit goods, have proliferated on (and disappeared from) the dark web since the introduction of Bitcoin [104, 199, 200, 201]. The amount of transactions in online marketplaces is vast and growing. For example, in 2020 Amazon reported a net revenue of 386B USD [202], while in 2019 the ecosystem of dark web marketplaces (DWMs) had reached a total volume of 4B USD [199].

Online marketplaces are commercial websites that allow participating buyers and sellers to exchange information about prices and products and to execute transactions [203, 204, 205]. Sellers can usually post an ad for their product that includes a product description, a price and a shipping method. Buyers instead can see all relevant product ads matching search keywords, and have access to reviews and seller ratings. When a purchase is made, the payment is processed through the platform, while shipping is usually taken care of by the seller.

Despite the importance of online marketplaces for e-commerce and global trade [206, 207], little is known about their empirical properties, transaction patterns and the resulting buyer-seller networks. The properties of the transaction network could, however, provide important insights into the presence of market power [208, 209], the nature of inter-platform competition [210, 211], product design [212], the effects of reputation on sellers' revenue growth [213], and the long-run sustainability of the platforms [214]. Moreover, measuring properties of the buyer-seller networks could help provide empirical foundations for theoretical models of online marketplaces, from the estimation of model parameters to suggesting specific model mechanisms. However, buyer-seller networks in online marketplaces have specific features

that make them different from other networks (e.g., social networks): they exhibit a naturally bipartite structure; most transactions (links) occur between anonymous agents; transaction activity might be infrequent and sporadic. Moreover, the structure of buyer-seller networks could depend on the nature of the traded products, on the types of buyers and sellers, on the user experience on the marketplace, or even on the legal, institutional and geographic constraints.

One strand of prior work relevant to this chapter has touched on various aspects of regulated online marketplaces. For example, the role of reputation and feedback [215, 216, 217] has been identified as one of the main drivers of the worldwide success of regulated online platforms [218]. Other work has looked at consumer search and the effect of rankings on product choice [219, 220, 221, 222], online auction markets [223, 224, 225, 226], market microstructure [227, 228] and price formation in online markets [229, 230, 231, 232, 233] (For a more complete but less up-to-date review see [203]). Another strand of research has studied unregulated marketplaces, as already discussed in chapter 2. This work has focused on country-specific studies [97, 98, 234], the effects of closures and law enforcement raids [102, 200, 199, 101, 198], the characterization of the trade of specific goods [235, 106, 236], the importance of geography [107, 237], or sociological interview-based studies [238, 234]. However, most work on unregulated online marketplaces was limited to specific markets, and focused on information available from public listings (e.g., using crawled data) [102, 107, 235, 104, 106, 236].

In this chapter, we focus on patterns in transactions which typically cannot be publicly observed either on regulated or unregulated online marketplaces. We analyze two datasets. The first dataset contains 220M transactions between 99M buyers and 7.4M sellers which occurred in 144 randomly sampled product markets of one regulated e-commerce platform between 2010 and 2020, for a total volume of over 10B USD. The second dataset contains 25M transactions involving 17M entities with a total volume of 4.2B USD which occurred in 28 major DWMs between 2011 and 2021, for a total volume of 4.2B USD (for more details on the datasets see *Data*). In both cases, the datasets cover all transactions which occurred in each corresponding market.

We observe striking similarities in user behaviour across online marketplaces, despite their significant differences. In section 5.2 we show how the number of transactions, amount, inter-event time and time between first and last transaction are highly heterogeneous across users but follow consistent fat-tailed distributions across all marketplaces. Then, we also show that individual behaviour is influenced by past purchases similarly (albeit less strongly) to what is observed in the renewal of past ties in social networks [239, 240]. Finally, in section 5.2.2 we propose a simple model of buyer-seller interactions that reproduces the main stylized facts of the data and emphasizes the critical role of preferential attachment [241, 242] and memory in the market dynamics. This chapter is based on publication [III]

## 5.1 Data

### 5.1.1 Dark Web Marketplaces

Our dataset contains the entire Bitcoin transaction history of 28 entities corresponding to DWMs between June, 2011, and February 2021 (see Fig. C.11 and Tab. C.2), as identified and preprocessed by Chainalysis (see chap 2). These markets all have an average daily volume of more than 15,000 USD, in order to be able to reliably measure different observables, and include all relevant DWMs as identified by law enforcement agencies [243]. The data contain all transactions received or sent by DWMs, excluding services such as exchanges (Bitcoin trading exchanges allow users to trade Bitcoin). Note that the data hide the direct buyer-seller link, because the money pass through the platform during the transaction.

We collect additional information on the analysed marketplaces from different sources, including the Gwern archive [244], law enforcement agencies reports and dedicated online forums. We focused our attention on the creation and closure dates of these markets, in order to correctly interpret the transaction data. We report the lifetimes of the selected markets in Fig. C.11, color coding by the daily average number of transactions as proxy for the market size. For more details on the markets, see also Tab. C.2.

### 5.1.2 E-Commerce Platform

The data used in this study consist in all the purchases made on 144 product markets from a popular e-commerce platform since 2010. The 144 product markets have been randomly selected from the markets that were active during the entire time period. The data cover only one geographical region. Similarly to the DWM data, the transaction data include: timestamp of the transaction, pseudonyms for buyer and seller, and the amount spent in the transaction. One key difference is that the data show the direct link between buyer and seller, forming a bipartite buyer-seller network and allowing for a more fine-grained analysis. For more details on the data see Fig. C.12.

## 5.2 Results

### 5.2.1 Empirical properties of buyer-seller networks

In order to characterize the buyer-seller networks, we start by analyzing different aggregate user-level quantities. First, we study the distributions (for each market) for the number and amount of user transactions. Results for all DWMs and for each e-commerce market are shown in Fig. 5.1a-d, where black and yellow lines are obtained by aggregating all users in the respective datasets. Single distributions display common behaviour, spanning several orders of magnitudes. It is important to note that distributions are computed without any rescaling or manipulation of the data, and that higher values generally reached by the regulated platform in all distributions are exclusively due to the different platform sizes. The

slight discrepancy between the distributions in the total number of received transactions can be ascribed to the different nature of the two datasets: while in the DWM dataset sellers can withdraw the earnings from several market trades at once, in the e-commerce data each transaction corresponds to a single purchase.

We then analyze the temporal dimension of the data. We focus on the distribution of user lifetimes, defined as the time between the first and last user transaction in the market, and the inter-event time between two successive transactions of the same user. Again, Fig. 5.1e-h shows remarkable regularities across different DWMs and different regulated product markets. In these distributions, as before, we also observe the effects of different sizes of marketplaces. The similarity between different distributions is particularly pronounced in the meaningful timescales between an hour and a month/year. Discrepancies for longer periods are due to the different lifetimes of the markets, whereas discrepancies for shorter timescales can be explained by the different nature of the two datasets: precise timestamps on transaction data for the regulated marketplaces vs. times at which the transaction is actually added to the Bitcoin blockchain (which depends on its algorithm) for the DWM dataset.




Figure 5.1: **Online marketplaces show strikingly similar patterns according to different aggregated quantities.** Top (a to d): Distributions of 4 analyzed users aggregate quantities: money and number of transactions both sent and received. Bottom (e to h): Distributions of 4 analyzed users temporal quantities: the inter-event time (time between successive transactions) and the lifetime (time between first and last transaction) both measured in days. Each blue line represents one DWM, the black line is the distribution built aggregating all DWMs together, each dashed red line represents one e-commerce platform market, while the dashed yellow line is the distribution aggregating all e-commerce markets. Similar patterns are observed between different markets in the same platform, but also across regulated and unregulated online marketplaces.

Having considered buyers and sellers separately, we now investigate the dynamics of buyer-seller relationships and the evolution of the buyer-seller network. We limit this analysis to e-commerce markets, since DWMs data do not contain buyer-seller links (see *Data* and chapter 2 for more details). We first consider how single users distribute their purchases across sellers: for example, buyers could purchase equally from multiple sellers or, alternatively, buyers could show loyalty to one seller from which they do most of their purchases. A standard way to quantify how distributed or concentrated this pattern is to compute the normalized entropy for the purchases of each buyer  $i$  as in Eq. 5.1, and then compute its distribution for all markets. The normalized entropy is defined as

$$e_i = - \sum_{j=1}^J n_i^j \log_2(n_i^j) / \log_2(J) \quad (5.1)$$

where  $n_i^j$  is the share of buyer  $i$ 's purchases from seller  $j$  and we sum over the  $J$  sellers the buyer purchased from. Fig. 5.2a shows that the distributions, computed for each market, are fat-tailed, with buyers populating the full  $[0, 1]$  support but with most of the mass towards the top, meaning that most buyers buy a similar number of times from the different sellers they purchase from. Buyers with zero entropy, who buy from just one seller, were excluded from the figure for visual clarity, but these were almost exclusively buyers who only made a single purchase (see Fig. C.1). In Fig. C.2 we further compare the distributions against a null model obtained reshuffling the transactions in the dataset (preserving buyers activity and sellers attractiveness), where we find lower heterogeneity and higher tendency towards high values of entropy. This implies that the empirical entropy distributions show a broader non trivial range of behaviours interpolating between perfect exploration and exploitation.

The observed normalized entropy distributions are compatible with different kinds of temporal patterns produced by two possible choices: either buyers choose to engage with new sellers they have never purchased from (i.e., exploration) or they return to a known seller (i.e., exploitation). We investigate these dynamics by leveraging insight from the social networks literature, where several studies have investigated how users explore and exploit social connections by renewing previously activated ties or by establishing new ones [239, 240]. Indeed, across different types of social networks, the temporal evolution of links that a person forms with their contacts can be captured by the following expression:

$$P(n) = (1 + n/c(k_{min}))^{-\beta(k_{min})} \quad (5.2)$$

where— now using the language of online marketplaces — $P(n)$  is the probability that a buyer (node) of degree  $n$  (who has already bought from  $n$  different sellers) chooses to buy from a new seller, while  $c$  and  $\beta$  are positive constants, depending on the final degree of the buyer, which measure their propensity to explore new sellers and thus the effect of memory. Following the procedure proposed in [239] (see section C.1 for more details), we group nodes in different classes according to their final degree: a buyer is in class  $k_{min}$  if the final degree  $k$




Figure 5.2: **Buyer memory affects their purchase decisions.** a) Normalised buyer entropy distribution for each e-commerce market (red), and the whole e-commerce platform (yellow), excluding users with zero entropy (mostly users with one transaction, see Fig. C.1) for visual clarity. The fat-tailed distributions span the full  $[0, 1]$  range, with most buyers almost equally buying from multiple sellers. b)  $P(n)$  is the probability to buy from a new seller after a buyer has already bought from  $n$  different ones. Each degree class  $k_{min} \leq k \leq 2 \cdot k_{min} - 1$  is rescaled according to the fitted value of  $c$  and  $\beta$  (see Tab. C.1 for the values), with Eq.2 (dashed line) well reproducing the memory effect on the buyers' behaviour: the more sellers they try, the less likely they are to buy from a new one.

satisfies  $k_{min} \leq k \leq 2 \cdot k_{min} - 1$ , starting from  $k_{min} = 2$ . For this computation, we aggregate all markets together in order to have a representative sample in classes with higher  $k_{min}$ . If a user is present in multiple markets, we keep its activity in different markets separated (i.e., effectively considering her as different users). We then fit Eq. 5.2 to each node class obtaining a value of  $c(k_{min})$  and  $\beta(k_{min})$  (see Tab. C.1).

Results are shown in Fig. 5.2b. Since different classes feature different values of  $\beta$  and  $c$ , we plot a rescaled  $P(x)^{1/\beta}$  as a function of  $n/c$ . Indeed, Eq.2 becomes  $1/(1 + x)$  (dashed line in Fig. 5.2b) for every degree class  $k_{min}$ , where  $x = n/c$ . In other words, we re-scale both axes assuming the empirical behaviour is captured by Eq.2. As shown in Tab. C.1, the parameter values are independent of the degree class and suggest a weaker ( $\beta \sim 10^{-1}$ ) effect than previously observed in social networks ( $0.48 \leq \beta \leq 2$ ) [240]. The close fit of the data to the predicted memory for different  $k_{min}$  indicates the applicability of Eq.2 in the dynamics of buyer-seller relationships. While users have different propensities to explore new sellers, they follow the same mechanism: the more sellers a user has bought from, the less likely is their next purchase from a new seller.

## 5.2.2 Modeling buyer-seller networks

In order to understand possible mechanisms that drive the properties of buyer-seller networks, we propose an agent-based model aimed at capturing the patterns observed in the previous section. The main features of the model are:

1. **Activity.** The rate at which buyers make transactions. As shown in Fig. 5.1, in both e-commerce and DWMs buyers feature heterogeneous propensities to make purchases.
2. **Memory.** When making new transactions buyers can either choose a seller they already bought from or pick a new one. As shown in Fig. 5.2b and discussed above, buyers have a memory of the sellers they had interacted with, and this memory affects their future purchases.
3. **Preferential attachment.** The attractiveness (i.e., popularity) of a seller is proportional to the number of their sales. This attractiveness captures the fact that, in online marketplaces, sellers are rated based on the feedback they receive from the buyers (i.e., customer reviews), and buyers prefer sellers with higher ratings, other things equal [216, 215, 245, 246, 247]. Here, we focus on the number of sales rather than sale volume to capture the fact that it is mainly frequency of transactions that matters for seller reputation. We don't use review scores, or other similar proxies, to measure the sellers' attractiveness as this data is not available in our dataset.

Given these three ingredients the model dynamics is as follows. The system consists of  $N$  buyers and  $M$  sellers. At  $t = 0$  we assign the activity  $a_i$  to each buyer  $i$ . Each seller  $j$  starts with attractiveness  $A_j = 1$ . At each time step  $t$ , each buyer makes a purchase with probability  $a_i \cdot \Delta_t$ , where  $\Delta_t$  is the simulation time step (fixed to 1 from now on). A buyer who interacted with  $n$  sellers in the past has probability  $P(n) = (1 + n/c)^{-\beta}$  of choosing a new seller and  $1 - P(n)$  of returning to a known one. In the first case, the buyer selects a new seller  $j$  proportionally to their attractiveness [248]  $A_j$ , in the latter, the buyer selects it proportionally to the number of previous interactions. In other words, buyers select sellers either according to past purchases or to their popularity. In both cases the attractiveness of the seller is increased by  $\mu$ . This model produces a bipartite temporal network: at each time step  $t$  we build a network in which two types of nodes— buyers and sellers—are linked if the buyer has purchased from that seller at time  $t$ . These networks are then combined together in an aggregated network, where each buyer-seller link is weighed according to the number of purchases between that buyer and that seller across time.

Compared to other activity-driven models developed to capture the temporal evolution of different social networks [239, 240, 25], our model extends the framework to bipartite networks of buyers and sellers and introduces the preferential attachment guiding the buyer selection process. Henceforth, we will refer to the model lacking preferential attachment, proposed in [239], as *Model NoPA*. We will also consider a version of the model that does not include the memory element (*Model NoMem*). Comparing these versions of the model will allow us to identify the role played by the different mechanisms.

A standard way to define and measure user activity in a (social) network is  $a_i = n_i / \sum_{\ell} n_{\ell}$ , where  $n_i$  is the number of purchases made by buyer  $i$ , where the sum is over all buyers in their market. In Fig. 5.3 we show the activity distributions of all e-commerce markets (a) and all DWMs (b). While curves exhibit fat-tailed behaviour, they no longer overlap due to different activity ranges and shapes in different product markets. As a result, we need to use market-specific empirical activity distributions as inputs for our model.




Figure 5.3: **Empirical activity distributions.** a): Activity distribution for all e-commerce markets in red, activity distribution of all aggregated markets in yellow. b): Activity distribution for all DWMs in blue, activity distribution of all aggregated DWMs in black. This plot is purely descriptive as the DWM activity distributions are not used in the rest of the analysis.

We now fit the model to the e-commerce data. As mentioned above, since the DWM dataset does not contain the full bipartite buyer-seller network, we cannot test all the model predictions on the DWM data. We employ a data-driven approach, fine-tuning the model to each single market so we can more faithfully compare the simulation results with the empirical buyer-seller networks (see section C.4 for more details). In the main text we show results for two different product markets, 26 more are shown in Fig. C.3-C.6, for a total of 28 markets (see section C.2 for the sampling procedure). We fix parameters  $\beta = 0.1$  and  $c = 0.001$  which we fitted previously (see Tab. C.1), and use the empirical activity distributions as measured in the data (see Fig. 5.3a) to reflect the observed heterogeneity across different markets. The value of  $\mu$  is determined with Maximum Likelihood Estimation for each market (see section C.3 for more details).

Results are in Fig. 5.4. We first compare the model’s output with the empirical distributions of the final seller attractiveness and degree. The attractiveness of a seller  $j$  is their market share  $A_j = s_j / \sum_{\ell} s_{\ell}$ , where  $s_j$  is the total number of sales of seller  $j$  and the sum is over all the sellers. Fig. 5.4 shows that the model reproduces both distributions well, while the NoPA variation of our main model (without preferential attachment) fails to capture the heterogeneity (up to six orders of magnitude) of these curves, emphasizing how preferential attachment is key to reproducing the presence of very active sellers. We then consider the buyer side of the network. We first study the degree distribution. Fig. 5.4 shows that the model captures the empirical distributions, while the absence of buyer memory generally leads to a small overestimation of the tails, since it does not induce the repetition of past interactions with a subset of buyers.




Figure 5.4: **The model reproduces different properties of buyer-seller networks.** Each row corresponds to a different market (see Fig. C.3-C.6 for other markets), whose simulation parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships.

Thus far, we have considered node-level properties aggregating detailed information on the links. For example, the attractiveness only accounted for the total number of links, whereas the degree only captures the total number of different buyers or sellers that the user has interacted with. To better understand how the model performs in reproducing finer details of the buyer-seller network structure, we test our model against two other properties of the aggregated network: link weight—the number of transactions between a buyer and a seller—and the buyer entropy, as defined in Eq.1. Our main model outperforms its two variations in reproducing the shape and tails of the link weight distribution. In particular, the memory mechanism appears to be fundamental in reproducing repeated transactions between a buyer and a seller. The buyer entropy distribution is again well-captured by the model and shows how the memory mechanism is key to capturing the diversity of relationships buyers establish with different sellers. Indeed, the *NoMem* model produces only entropy values close to 0 and 1; this happens because without memory, a buyer almost never finds any previous seller, hence buyers making more than one purchase almost always buy from new sellers.

We have seen that our model is able to capture various aspects of the final aggregated buyer-seller network. The next step is to see whether our model can also reproduce the temporal evolution of the buyer-seller network. To investigate this, we focus on the degree of top sellers since we previously showed these sellers generate the largest activity and volume in these markets. We measure time by the total number of purchases made. Results

are shown in Fig. 5.5a-c, where we plot the temporal evolution of the top 50 (a), 100 (b) and 200 (c) sellers degree distribution for one illustrative product market. In doing so, we compare the model to its two variations and the data. Results for more product markets are shown in Fig. C.7-C.10. The main model is able to reproduce the temporal evolution of the distributions, as clearly shown by the cores (i.e., interquartile ranges) overlapping at different times. We further compute the absolute value of the difference between the mean of the models' distribution and the mean of the data, for each of the nine equally spaced time steps and for all 28 simulated product markets. As shown in Fig. 5.5d-f, the model is better able to reproduce the temporal dynamics for all simulated markets. Indeed, the median of the distance distributions is always smaller in the main model than the two other model variations.




Figure 5.5: **Model reproduces the temporal evolution of the top sellers degree distribution.** Top (a to c): temporal evolution of the degree distribution of the top 50 (a), 100 (b) and 200(c) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5<sup>th</sup> to 97.5<sup>th</sup> percentiles. Results are shown for one product market, all other markets are shown in Fig. C.7-C.10. Bottom (d to f): Distribution of the distance between the empirical and model(s) median degree of the top 50(left), 100(center) and 200(right) sellers, for all product markets and time steps, and the three considered models. Vertical lines represent the distributions median, showing that the model median is always smaller than the alternatives. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism.

## 5.3 Discussion

In this chapter, we have analyzed 244M (25B USD) transactions occurring on regulated and unregulated online marketplaces. First, we have revealed remarkable regularities in the aggregate static and temporal properties of the buyer-seller networks, both for buyers and sellers. Then, we have revealed how buyers are affected by the memory of past interactions. Finally, we have proposed a model which captures the main stylized facts of the data, based only on three well-known network formation mechanisms of online marketplaces: buyers have different propensity to make purchases, they remember the sellers they purchased from, and they are more likely to buy from successful sellers.

It is important to highlight the limitations of our study, which also represent directions for further work. First, while our study is based on (pre-processed) blockchain data, access to DWM server logs could provide more detailed information on some specific markets, for instance, the directed buyer-seller links which are not observable in our data. Second, the model could be further enriched with other known mechanisms: pricing dynamics [249, 250], product search ranking [219, 220, 221, 222], variable (e.g., also negative) customer reviews [217], sellers entering or leaving the platform [251], and recommendation algorithms [252]. Finally, including richer economic incentives (e.g., strategic behaviour) to model buyers' and sellers' decisions could shed light on how agents could exploit their market power. In particular, the inclusion of strategic behaviour would also drop phenomenological rules such as preferential attachment, which would naturally result from the agents' behaviour [253, 254]. A deeper understanding of economic incentives and equilibrium behaviour in buyer-seller networks could ultimately inform market design and regulation of online marketplaces.

Nevertheless, our work supports and extends previous findings. The fat-tailed heterogeneous curves in Fig. 5.1a-d substantiate previous observations of high concentration in DWMs: wholesale [237], few sellers [102] or few buyers [107] were found responsible for the largest part of volumes in smaller samples of data. The fat-tailed inter-event time distributions, spanning times between a second and a year, are compatible with the bursty nature of several social activities [255, 256], and the finding about a shared memory kernel further points to a similarity between social and economic activities [239, 240]. Taken together, our results could inform and enrich economic models where heterogeneity assumptions are now commonplace [211] but empirical evidence on the structure of buyer-seller networks has not yet been introduced.

The regularities observed in Fig. 5.1 are surprising given the differences in the marketplaces covered by our data: transactions on the clear web with state enforcement of contracts [257] vs. transactions on the dark web that rely mainly on reputation and self-governance [246]; the sale of only regulated products vs. mainly unregulated products; the use of fiat vs. the use of cryptocurrencies. And, indeed, there is both substantial heterogeneity in product markets in the e-commerce dataset and several differences across marketplaces in the DWM dataset (e.g., existence time period, geography, product focus, etc.). Our model suggests specific mechanisms that drive the regularities across the two datasets. Sellers build

a reputation that makes them more attractive to buyers who, in turn, are affected by their memory of the sellers they already purchased from. In particular, the presence of both memory and preferential attachment is fundamental in reproducing both local and global properties of the buyer-seller network, as already shown for the intrinsically different social networks [241, 242, 255, 239]. However, commercial interactions exhibit important differences compared to social interactions, with preferential attachment playing a dominant role in the market dynamics.

Our results point towards alternative strategies to attempt to reduce trading of illicit goods on dark web marketplaces. Historically, DWMs have been closed after long and expensive operations targeting the market admins in order to arrest them and shut down the servers [243]. However, the high degree of concentration, the importance of preferential attachment, and the memory kernel in the buyer dynamics, all suggest that limited observations of the market dynamics could give a clear enough picture of who the key actors of these networks are. For instance, key sellers will most likely attract most of the observed purchases from the more active buyers, and stopping them would effectively stop a large part of the market trade. In this regard, our model could also be used to produce candidate synthetic DWM buyer-seller networks to quantitatively study and simulate the effects of targeting “key players” on marketplaces [258].

Finally, a better understanding of buyer-seller network formation could have consequences for market design and regulation. For example, fat-tailed distributions show a high degree of concentration on both buyer and seller sides of the marketplaces: just a few agents (both on the buyers and seller sides) are responsible for a vast majority of the transaction volume. While buyer market power appeared in analyzes of labor monopsony and retailers [259, 260], our empirical finding of buyer concentration calls for a deeper understanding of buyer power in online marketplaces. Moreover, these observations can also inform theoretical economic models of online marketplaces, providing empirical backing to heterogeneity assumptions and suggesting specific values for parameters or shapes for distributions. Also, we find signs of both local (memory) and global (reputation) mechanisms in the structure and evolution of buyer-seller relationships. Thus, the inclusion of memory and reputation in previously developed models can improve our understanding of the pricing of network effects [209], inter-platform competition [211] and long-run sustainability of the platforms [214].

## Chapter 6

# The decentralized evolution of decentralization across fields: from Governance to Blockchain

“For students of recent domestic affairs it is becoming increasingly evident that ‘*decentralization*’ is a magic word”. With these words in 1975 Herbert London opens his article “The meaning of decentralization” [261]. Almost 50 years later, Schneider states that *decentralization* “is called for far more than it is theorized or consistently defined” [262]. *(De)centralization* (i.e., either *Decentralization* or its counterpart *Centralization*) has indeed become almost a buzzword, permeating not only the academic literature, but also the public discussion. The debate between centralized and decentralized contact tracing at the beginning of the COVID-19 pandemic is a clear example [263]. However, one of the major drivers of its growth has certainly been the rise of blockchain based technologies such as cryptocurrencies, NFTs and the metaverse. [3, 264, 265, 266].

However, *(de)centralization* is not a new concept, and has different connotations across fields. In political science, it usually refers to the delegation of power to local communities with respect to a central government [267]. The concept has similar connotations when referring to educational [268], fiscal [269] and more generally governance systems. Other domains where the term is widely used include public health [270], internet protocols [271], robot swarms [272] and social network analysis [273] among others, with the last one providing one of the few quantitative definitions available thanks to Freeman in 1978 [274]. Given this background, some questions naturally arise: have these different disciplines independently developed the concept of *(de)centralization*, maybe even with different meanings (i.e., a case of polysemy)? Have they influenced each other? Which fields have been most influential to the evolution of this concept?

Here, we address these questions by studying a corpus of scientific literature indexed by the Semantic Scholar open database [20]. In section 6.1 we describe the data science

pipeline we developed, and publicly released, for this chapter. First, we describe the main data source: S2AG, and how to collect and preprocess the data. Second, we describe the hierarchical Stochastic Block Model, and how one can employ it in our case to cluster both articles into fields, and title words into topics, to gain a complete picture of the selected subset of the literature. Finally, we describe the methodology of knowledge flows, which allows us to compute significant flows of knowledge, through paper citations, flowing between different clusters (fields) obtained with the previous method.

In section 6.2 we describe the main results obtained in this chapter. In section 6.2.1 we employ this pipeline to study the concept of decentralization. First, we observe an exponentially growing interest in the topic, with an author in 154 contributing to a paper mentioning *(de)centralization* in its title or abstract in 2021. Then, we map the literature on *(de)centralization* by focusing on the subset of relevant articles and clustering them according to their semantic and citation information. This way, we discover that different academic fields have separately contributed to this topic. We hence study how the different clusters have influenced each other, showing how much more transfer of knowledge between different academic areas is happening in recent years. Interestingly, our analysis reveals that STEM and social sciences did not influence each other. Finally, in section 6.2.2 we focus on two paradigmatic examples: Governance, interpreted generally as “the way that organizations or countries are managed at the highest level, and the systems for doing this” [275], and Blockchain, including all blockchain based technologies, from cryptocurrencies to NFTs and the metaverse. We show how Governance is the first cluster to extensively make use of the term *(de)centralization*, containing the most or second most number of papers each year since its appearance in the 1950s, and playing a leading role in the transfer of knowledge to other fields until the 1990s. Blockchain instead has become both the most influential cluster and the most productive cluster in the past 10 years, showing three different phases in its recent history characterized by different interactions with other fields. Overall, our results shed light on the history and evolution of the more and more important concept of *(de)centralization*. Furthermore, we publicly release the code of the pipeline developed for this chapter, so that it may be used to study and understand the evolution of other concepts through the lenses of the academic literature. This chapter is based on publication [IV].

## 6.1 The pipeline

In this section, we briefly describe the pipeline we have set up and publicly released<sup>a</sup> to select the data and perform the research described in this study. The pipeline is conceptually divided into three steps: (1) data collection, (2) clustering of the dataset using a multilayer hierarchical stochastic block model, and (3) analysis of the influence between clusters over time using knowledge flows.

---

<sup>a</sup>See <https://github.com/alberto-bracci/decentralization>.

### 6.1.1 Step 1: Data Collection

The first step consists in collecting the academic publications related to the concept of *(de)centralization*, or potentially to other concepts. To perform a large scale analysis of the academic literature, we exploit the possibility to access the publicly available *Semantic Scholar Academic Graph* (S2AG, pronounced “stag”), which provides monthly snapshots of research papers published in all fields [20]. Launched in 2015 by the Allen Institute for Artificial Intelligence (AI2), Semantic Scholar provides this corpus as an open access database with the specific scope of facilitating scientific analysis of academic publications. It contains about 203.6M papers (1<sup>st</sup> Jan. 2022 snapshot), 76.4M authors, and 2B citations. Moreover, this database recently incorporated the Microsoft Academic Graph (MAG) [276], which was shut down at the end of 2021 [277].

From this corpus we extracted the data about papers that contain the root string “*centraliz*” or “*centralis*” in words of the title or abstract, to capture possibly all variations of words related to the concepts of *Centralization* and *Decentralization* (nouns, adjectives, etc.). In this way, we also incidentally captured articles written in different languages, mainly Portuguese, French and Spanish, and also a minority of unrelated articles (e.g., biology articles involving plant species including “*centralis*” in their name). The resulting dataset has 425k papers characterized by a series of attributes. Among these, of particular interest to us there is the title, the abstract, the authors, all in- and out-citations (respectively citations and references), the year of publication, and the fields of study, which were determined based on machine learning field classifiers leveraging on the existing MAG taxonomy and classification [278]. Notice, however, that some articles miss one or more of these attributes. See Table D.1 for details on how many papers have each of these attributes.

### 6.1.2 Step 2: Hierarchical clustering

In the Semantic Scholar corpus almost each paper is associated to a list of fields of study. However, these are high-level, as there are in fact only 19 fields of very heterogeneous sizes (see Table D.2 for details on how many papers are classified in each field of study). Moreover, sometimes the fields are not correctly assigned. In the second step of the pipeline, we hence use a multilayer hierarchical stochastic block model (hSBM) [26, 279], developed to find statistically significant clusters at multiple hierarchical levels for the analysis of text data with multiple data types. Here, in fact, we consider two layers. The document layer—where links represent citations between papers—and the text layer—a bipartite network between documents and the words present in their titles. The two layers are connected since the document nodes are the same on both layers. The method naturally produces clusters of documents and topics (word clusters), incorporating the information from both layers in the process. Furthermore, as the name suggests, the model produces a hierarchical clustering, providing a richer structure of both article clusters and topics, which captures both small clusters and topics and how they are related to each other in a higher level structure.

We consider only the papers in our dataset that have a non-empty title and contain at

least one citation or reference to another paper in the dataset (42.7% of the initial dataset), as we are interested in how the concept of *(de)centralization* evolved in the academic literature, and citations are the most natural proxy for how knowledge is transferred. We use title texts, instead of abstracts, for various reasons: firstly, because the title is more frequently available than the abstract (see Table D.1); secondly, because the title has the advantage of being more distilled compared to abstracts [280]; lastly, because titles contain a significantly smaller number of words than abstracts, allowing us to obtain a text layer similar to the document layer in terms of number of edges by simply cutting out words present in less than 5 documents. It is indeed well known that the hSBM performs optimally when both layers have roughly the same size, otherwise the smaller layer is effectively ignored by the algorithm [279]. The filtered dataset hence consists of 181,605 documents and 15,381 different words, summing up to 590,215 document-to-document citation links and 1,396,830 document-to-word links.

To make sure results are robust, the algorithm is run 100 times, and the consensus partition between the 100 runs is then computed. Afterwards, keywords are assigned to each cluster to roughly represent the content and themes of the articles within them (for more details see Fig. D.2, Fig. D.3 and Table D.3, with related section). Keywords are chosen by looking at the most frequent words in the cluster, the most significant topics in the cluster according to the normalized mixture proportion [279], as well as the first 5 papers in the cluster according to different measures (see Section D.2.1 for more details).

### 6.1.3 Step 3: Knowledge flows

In the third step of the pipeline, we want to better understand how the different groups of documents identified by the hSBM have influenced each other throughout history. To do so, we evaluate the knowledge flows between these groups, using article citations as proxy [27]. In particular, we compute the knowledge flow from one cluster  $a$  in one year  $Y_a$  to another cluster  $b$  in a future year  $Y_b$ . The computation takes into account the fraction of citations towards papers in  $a$  of the year  $Y_a$  from papers in  $b$  published in the year  $Y_b$  with respect to the fraction of citations towards  $a$  in  $Y_a$  from all papers published in  $Y_b$ , as well as the overall fraction of papers of  $a$  in  $Y_a$ . The citation network suffers indeed from a series of inherent biases: field size, typical number of citations in a field or a journal, typical number of references, age of the fields etc. This method de facto considers the number of citations with respect to a null model, resulting in a link weight which is effectively a z-score.

Mathematically, if a collection of papers is divided in a partition  $\mathcal{P}$  of clusters such that different clusters do not overlap and altogether form the collection of papers, then we can define the knowledge flow units  $C_{a \rightarrow b}(Y_a, Y_b)$  from papers in cluster  $a \in \mathcal{P}$  published in the year  $Y_a$  to papers in cluster  $b \in \mathcal{P}$  in a future year  $Y_b$  by counting how many citations have occurred from  $b$  to  $a$  in the two years, that is,

$$C_{a \rightarrow b}(Y_a, Y_b) = |\{(x, y) : x \in a, Y_x = Y_a, y \in b, Y_y = Y_b \text{ s.t. } y \text{ cites } x\}|. \quad (6.1)$$

As said before, we need to normalize this number with respect to a null model, so as to keep into account different sizes of clusters and different norms in citation practices. Hence,

the knowledge flow  $K_{a \rightarrow b}(Y_a, Y_b)$  from  $a$  in year  $Y_a$  to  $b$  in year  $Y_b$  can be computed in the following way:

$$K_{a \rightarrow b}(Y_a, Y_b) = \begin{cases} 1 & \text{if } \frac{C_{a \rightarrow b}(Y_a, Y_b)}{\sum_{c \in \mathcal{P}} C_{c \rightarrow b}(Y_a, Y_b)} \Big/ \frac{|x \in a : Y_x = Y_a|}{\sum_{c \in \mathcal{P}} |x \in c : Y_x = Y_a|} \geq 1 \\ 0 & \text{otherwise} \end{cases}, \quad (6.2)$$

After the normalization against the null model, knowledge flows can be indeed treated as z-scores. Hence, in Eq. (6.2) we consider a knowledge flow as significant (i.e., a binary value of 1) if higher than the threshold 1, and as not significant (i.e., 0) otherwise.

Therefore, we obtain a binary value for each pair of clusters and each pair of years. In other words, the collection of (knowledge flow) links between all pairs of clusters and years generates a temporal network of clusters, which we aggregate in different ways to facilitate the following analysis. In particular, we consider the average knowledge flow  $K_{a \rightarrow b}(Y_a)$  from a cluster  $a$  to another  $b$  from a specific year  $Y_a$  as the average of the knowledge flows  $K_{a \rightarrow b}(Y_a, Y_b)$  from cluster  $a$  to  $b$  from year  $Y_a$  to all years  $Y_b > Y_a$ , taking into account only years  $Y_b$  where there is at least one publication in  $b$ . Formally, this reads:

$$K_{a \rightarrow b}(Y_a) = K_{a \rightarrow b}(Y_a, \bullet) = \langle K_{a \rightarrow b}(Y_a, Y_b) \rangle_{\{Y_b > Y_a : \exists x \in b \text{ s.t. } Y_x = Y_b\}}. \quad (6.3)$$

This value represents, on a scale from 0 to 1, how much publications in cluster  $a$  in year  $Y_a$  have influenced the future of cluster  $b$ . Analogously, we define the average knowledge flow  $K_{a \rightarrow b}(T)$  from cluster  $a$  to cluster  $b$  from a period of time  $T$  to the future by averaging  $K_{a \rightarrow b}(Y_a)$  over all years  $Y_a$  in  $T$  in which there is at least one publication in  $a$ , that is,

$$K_{a \rightarrow b}(T) = \langle K_{a \rightarrow b}(Y_a) \rangle_{\{Y_a \in T : \exists x \in a \text{ s.t. } Y_x = Y_a\}}. \quad (6.4)$$

We can also measure the average influence in terms of knowledge flows from a cluster to all other clusters and vice-versa, as well as the average knowledge flow among all clusters, respectively as follows:

$$K_{a \rightarrow \bullet}(Y) = \langle K_{a \rightarrow b}(Y) \rangle_b, \quad K_{\bullet \rightarrow a}(Y) = \langle K_{b \rightarrow a}(Y) \rangle_b, \quad K_{\bullet \rightarrow \bullet}(Y) = \langle K_{a \rightarrow b}(Y) \rangle_{a,b}. \quad (6.5)$$

Here,  $K_{a \rightarrow \bullet}(Y)$  refers to the average influence from papers in cluster  $a$  published in year  $Y$  towards all clusters in the future. On the opposite,  $K_{\bullet \rightarrow a}(Y)$  refers to the average influence of the papers in all clusters in the year  $Y$  towards the future of cluster  $a$ . Finally,  $K_{\bullet \rightarrow \bullet}(Y)$  refers to the average influence (towards the future) of all papers published in year  $Y$ .

## 6.2 Results

### 6.2.1 The decentralized evolution of (de)centralization

We start by analysing the number of papers mentioning *(de)centralization* over the years (see *The pipeline* section for more details), comparing it to the total number of academic outputs

(papers, books etc.) produced in time, which is known to be exponentially increasing [281]. As shown in Fig. 6.1(a), the fraction of papers mentioning *(de)centralization* has been exponentially increasing in time since the 1950s, rising to around one paper in every 315 in 2021. The growing interest in this topic is also reflected by the increasing number of authors involved in such academic research. Indeed, as shown in Fig. 6.1(b), the fraction of authors producing such research has risen exponentially by more than one order of magnitude, with almost one academic every 154 writing a paper mentioning the topic in 2021. This growth is also seen in terms of raw number of publications and authors, as shown in Fig. D.1, where we compare these numbers for the S2AG corpus and the *(de)centralization* dataset and find a stronger exponential rise for the latter. Notice that for both papers and authors there are some periods with a higher or lower increase in the fraction, showing spikes of interest at particular times. For example, in Fig. 6.1(a,b) we can see that between the late 1970s and the 1980s the growth rate was faster than the overall exponential fit.

In order to understand what has characterized the origins and evolution of the topic, we set to identify topics and clusters of papers in the dataset by using the hSBM algorithm [26, 279] described in *The pipeline* section. In the following analysis, we focus our attention only to years after 1950. Before this date there are only around 100 papers in our *(de)centralization* dataset. The very first is a political science one from 1851 on local self-governments versus centralized governments [282]. Among the others in this period, apart from around 50 papers that relate to *(de)centralization* in governments, organizations and states, we have detected 30 papers that are actually false positives of the selection process. Considering also how, in general, digitalization issues may have contributed to the small number of papers before 1950, the reliability and coverage of the first 99 years of the data are unclear, and we opted to exclude them from the analysis.

The hSBM algorithm identified 7 hierarchical levels of clusters of documents. On the left of Fig. 6.1(c), we draw this hierarchy only until the 3<sup>rd</sup> level (starting from the common root at level 0) for visualization clarity. On the right, we also show the heatmap of the number of papers in time for each cluster at the 3<sup>rd</sup> level, for a total of 16 different clusters after excluding other 5 clusters with less than 500 papers not included in this analysis. The keywords shown in the heatmap have been manually assigned to roughly characterize each cluster. In particular, the keywords between parentheses have been chosen amongst the most frequent and most significant in the clusters at the 4<sup>th</sup> level, while the most representative keyword at the 3<sup>rd</sup> level has been chosen and printed in bold. For details on how they were assigned see *The pipeline* section and Sec. D.2.1. In the following, we refer to a cluster at the 3<sup>rd</sup> level by its representative keyword (capitalized). As shown by the hierarchy and by the horizontal red lines in the heatmap, clusters are divided in three main branches. Looking at the two biggest branches, we can see a clear division between more STEM oriented documents (top branch) and those in Political sciences, Social sciences, as well as Medicine (bottom branch). Notably, a third smaller branch appears isolated, including papers at the intersection of the other two, mostly about Wireless technologies and their applications. Going into more details, in the STEM branch we notice how Cybersecurity, Control theory




Figure 6.1: **The rising interest of academic literature towards (de)centralization.** (a) Fraction of papers mentioning (de)centralization in the Semantic Scholar corpus. (b) Fraction of authors mentioning (de)centralization in the Semantic Scholar corpus. Both fractions have been steadily increasing since the 1950s, showing growing interest in the topic. (c) Number of papers in the clusters at the 3<sup>rd</sup> level of the hierarchy in each year. Clusters are ordered respecting the hierarchical network on the left, in which node and link sizes are proportional to the total number of papers in the related cluster. In the heatmap, white lines individuate clusters belonging to the same cluster at the 2<sup>nd</sup> level of the hierarchy, while red lines divide different clusters at the 1<sup>st</sup> level. The representative keyword of the clusters at level 3 is reported in bold in the respective rows, while all the specific keywords identified at the 4<sup>th</sup> hierarchical level are shown within brackets. Clusters with less than 500 papers in total are not shown in the figure.

and Telecommunication (clusters 1, 6, and 10 respectively) are the ones producing most publications, with Blockchain (cluster 3) becoming the most relevant in the last 5 years in terms of number of papers published per year. On the other branch, Governance (cluster 12), including works in Political science, Education and Fiscal federalism, is the most relevant cluster, while Environment, Social network analysis and Health clusters (respectively clusters 13, 14, and 15) have produced a smaller number of papers. Furthermore, see Fig. D.4 and Fig. D.5 for a similar plot done at the 4<sup>th</sup> level and for a bipartite hierarchical network showing how clusters are represented in the various topics respectively.

As said, in Fig. 6.1(c) we show how the number of papers in each of these clusters has

evolved over time. Looking more into details of the early history of *(de)centralization*, the first papers adopting the term have all been in the Social sciences branch, most importantly the Governance cluster, followed by Social network analysis and Health. In the 1950s, indeed, there are 58 papers in Governance, which represents the first cluster to adopt the term and use it extensively. Some of these articles refer, among other things, to democracy as a form of centralized decision-making system [283]. Other clusters with more than 10 papers refer to Social network analysis and Health, as seen for example in Kaufman's "Toward an interactional conception of community" [284]. Here, *centralization* is depicted as a force gradually destroying the concept of community as a social unit. Notably, most of these papers have no citations from other articles in the full corpus, with only some citations within the cluster of governance. In the 1960s, the largest growth is found again in the Governance and Health clusters, both reaching around 150 papers each in the decade. An important example of the former is that of Bachrach et al. [285], where they highlight how different disciplines (i.e., social and political sciences) reach completely opposite conclusions about the *(de)centralization* of power. In the same decade *(de)centralization* also appears in other relevant clusters, namely Social network analysis (50 papers) and Investments (29 papers), with a significant number of citations in both directions between them. The term is picked up from the STEM branch only later in the 1970s, especially through works in Control theory and Optimization [286], coming significantly to a popular domain as Cybersecurity only in the 1980s.

We have seen how different domains have picked up the concept of *(de)centralization* at different times. It is therefore natural to ask whether they developed such uses separately, or they influenced each other in some way. The hierarchical clustering partially answers this question, as it gives a degree of separation between domains based on citation and semantic information. However, significant information is still present in the citations between papers of different clusters. We exploit this by computing knowledge flows [27], whose aim is to quantify the transfer of knowledge given by the citations between groups of papers through a comparison with a null model (for more details see *The pipeline* Section). Here, we study the average knowledge flow  $K_{a \rightarrow b}(T)$  from papers in a cluster  $a$ , at level 3 of the hierarchy, in a period of time  $T$  to future publications in another cluster  $b$ , represented by a number between 0 and 1 showing how significant this influence has been. In particular, in Fig. 6.2 we consider three different periods of time  $T$ : 1970-1989, 1990-2007 and 2008-2020. Similarly to what we will do in the next figures, We have excluded the year 2021 as a source of knowledge flow, because our dataset ends at the end of 2021, thus meaning that we cannot evaluate knowledge flows from papers of 2021 to future years. In the figure, all clusters are ordered as in Fig. 6.1(c), with the representative keywords shown in the legend below. For each period  $T$ , the color of the cell of row  $a$  and column  $b$  of the heatmap refers to the average knowledge flow  $K_{a \rightarrow b}(T)$  from papers in cluster  $a$  in that period of time to future papers in cluster  $b$ , according to the colormap shown below. As seen in Fig. 6.2(a), between 1970 and 1989 clusters have little to no influence on the future of the other ones. We have previously seen how in these years the use of *(de)centralization* started to rise across some domains, mostly being Governance, Control theory, Social network analysis, Health, Cyber-




Figure 6.2: **Temporal evolution of the influence between clusters.** Average knowledge flows  $K_{a \rightarrow b}(T)$  from each cluster  $a$  to each cluster  $b$  at level 3 in the period  $T = 1970-1989$  (a),  $T = 1990-2007$  (b) and  $T = 2008-2020$  (c), represented by the different colors according to the colorbar in the bottom left part of the figure. A representative keyword for each cluster is reported on the bottom right part of the figure. White lines denote clusters belonging to the same 2<sup>nd</sup> level cluster, whereas red lines mark different branches at the 1<sup>st</sup> level. In the first period, little to no communication is happening between different clusters. In recent years, more communication happens inside the same 2<sup>nd</sup> level cluster, and towards the middle branch (cluster 11). However, little communication happens between the two other branches, roughly representing the STEM and social sciences communities respectively.

security, and Investments. However, apart from Governance and Control theory (clusters 12 and 6), these clusters have low knowledge flow even to themselves, meaning that the use of *(de)centralization* was only relegated to sporadic and not so influential papers in the literature. This also confirms that the topic has appeared independently at this early stage. In Fig. 6.2(b) instead, we can see how much more transfer of knowledge has occurred between clusters from 1990 and 2007. As shown in Fig. 6.2(c), this trend is even more pronounced in recent years, which notably coincides with the creation and rise of blockchain technologies. Interestingly, these transfers reflect the structure of the hierarchy and denote significant differences between the high-level domains. The STEM branch (made of the clusters 0 to 10) shows clear communication between clusters belonging to the same group both at the 2<sup>nd</sup> and 1<sup>st</sup> level (respectively within white and red lines), whereas the right bottom branch shows almost no communication with the other domains, especially after 2008. The only significant knowledge flow from this branch in the middle period goes from Environment (cluster 13) to Optimization (cluster 5), while in the last period this is only relegated between Environment and Smart grids (cluster 2). The middle branch instead shows clear influence from the other two, and little influence towards them, especially in the last period. Moreover, notice how

the highest knowledge flows between different clusters in Fig. 6.2(b) are from those that, in the first period, were starting to be more influential, while in the last period it is relegated mostly to just STEM clusters. Put together, the heatmaps show a clear decentralized birth of the concept of *(de)centralization*, appearing in different fields and domains with little to no communication between each other. Instead, in recent years, we find a more coordinated evolution, even though still sectorial in some cases, and mainly led by STEM related clusters.

## 6.2.2 The case of Governance and Blockchain

Having analyzed the concept of *(de)centralization* in the general academic landscape, we now focus on two of the most important clusters in the history of this topic: Governance and Blockchain. As shown in Fig. 6.1(c) and in Fig. D.6, these two clusters are among the biggest across time in terms of number of papers. The Governance cluster has always been first or second with respect to the other clusters at the 3<sup>rd</sup> level, while Blockchain was barely present before 2008, the year of the Bitcoin white paper [3]. After that, Blockchain gradually increased in size and had an exponential explosion after 2015, coincidentally with the increasing hype around the technology and its applications, in particular Bitcoin and ethereum [287, 288, 289]. Finally, it has become the most productive cluster since 2019, surpassing governance.

To better understand their role in the evolution of the literature on *(de)centralization*, we consider the average knowledge flows between clusters for each year, that is looking at  $K_{a \rightarrow \bullet}(Y)$ ,  $K_{\bullet \rightarrow a}(Y)$ , and  $K_{\bullet \rightarrow \bullet}(Y)$ , defined in Eq. (6.5). Therefore, in Fig. 6.3 we rank clusters year by year using  $K_{a \rightarrow \bullet}(Y)$  in (a) and  $K_{\bullet \rightarrow a}(Y)$  in (b), i.e., looking at how much the papers of a cluster  $a$  in a year  $Y$  have influenced, on average, the future of all other clusters (a), or, vice versa, how much all clusters have influenced the future of  $a$  (b). From these plots we can see how, on the one hand, Governance has been in the top ranks until the late 1980s, both as a source and target of knowledge flows. However, in the early 1990s it started to decrease in importance, reaching the bottom ranks in the 2000s, despite being the first cluster in terms of number of papers each of these years. On the other hand, in Fig. 6.3(a) we notice that the rise of Blockchain started only in 2010, being almost always outside of the *(de)centralization* literature discussion until this point. Then, very sharply, Blockchain becomes the first cluster in terms of influence towards other clusters in 2013, maintaining its position in the following years. Hence, the literature on Blockchain has been key in the development of the *(de)centralization* discussion in the most recent years. Moreover, looking at Blockchain in Fig. 6.3(b), papers of other clusters before early 2000s have had almost no impact on the scientific future of Blockchain. Interestingly, it has received a lot of influence from publications between 2006 and 2012, that is about when the blockchain and Bitcoin originated [3], as well as after 2017, mostly due to the increasing amount of applications using blockchain in the most diverse contexts in recent years. Finally, notice the loss of influence on Blockchain from papers between 2013 and 2016.

These results are corroborated by the time evolution of the average knowledge flow com-

pared to the overall average  $K_{\bullet \rightarrow \bullet}(Y)$ . Indeed, in Fig. D.7 we show how Governance has been increasingly important in influencing other clusters until the 1980s, while since the 1990s it has had a lower average knowledge flow than the average among all clusters. Similarly to what is shown by the ranks, after 2013 Blockchain starts to have a much higher influence towards the other clusters compared to the average.

We have seen how influential Governance has been in the early literature about *(de)centralization*, and how Blockchain has risen in recent years as the most important influential cluster, contributing in terms of knowledge flow towards other branches of literature. It is therefore a natural next step to investigate in more details which clusters in particular have influenced or have been influenced by Governance first and Blockchain then, and see how these interactions have changed over time. We start this analysis from the more recent case of Blockchain. This cluster started to appear only around 2008 with the Bitcoin white paper [3]. Moreover, we notice a decrease in the influence on this cluster in mid 2010s. We therefore divide the 2008–2020 time span in three parts, following blockchain history: 2008–2014, representing the origin of blockchain applications before the advent of ethereum; 2015–2018, when the field got more recognition thanks to ethereum and Bitcoin; and the final 2019–2020 period, in which we have seen the explosion of academic literature production and the widespread success of multiple applications such as DeFi, NFTs and the metaverse.

In Fig. 6.4 we plot, in a decreasing order, which clusters have been most influenced by




Figure 6.3: **Ranking the influence of Blockchain and Governance in the *(de)centralization* literature.** (a) Ranking in time of the influence coming from a cluster in the 3<sup>rd</sup> hierarchical level, computed on the average knowledge flow  $K_{a \rightarrow \bullet}(Y)$  from papers in cluster  $a$  published in the year  $Y$  towards all other future papers. (b) Ranking in time of the influence to a cluster in the 3<sup>rd</sup> hierarchical level, computed on the average knowledge flow  $K_{\bullet \rightarrow a}(Y)$  from papers published in the year  $Y$  towards future papers in cluster  $a$ . The blockchain cluster, highlighted in blue, has become a central actor in the recent literature on *(de)centralization*, supplanting the governance cluster, highlighted in red.

(a) and have most influenced (b) Blockchain during the three periods. To this end, we use a Sankey diagram, showing how the overall picture has changed in the three different phases. The plot is done using the average knowledge flows  $K_{a \rightarrow b}(T)$ , where  $T$  is the selected period, while  $a$  and  $b$  are fixed to Blockchain in Fig. 6.4(a) and Fig. 6.4(b) respectively. We can see important differences across the three periods. First, as shown by Fig. 6.4(a), the early literature of Blockchain has had a big impact on most of the clusters. As a matter of fact, there are only a few cases where the average knowledge flow from Blockchain to another cluster is zero, shown by a circle in the respective node and a lighter color in the corresponding link. We also notice that Cybersecurity, Smart grids, Edge-computing, Wireless technologies, and Federated learning have a very significant average knowledge flow from Blockchain, i.e.,  $K_{a \rightarrow b}(T) > 0.1$ , shown by the double stars, while other clusters with  $0.01 < K_{a \rightarrow b}(T) \leq 0.1$  are represented with only one stars. Notice how Blockchain has continued to have a big impact on these mentioned clusters. In particular, papers of Blockchain in the last period have had a significant impact on the future of only Smart grids and Wireless technologies, as well as of Cybersecurity to a lesser extent. On the contrary, there is no significant knowledge flow to all other clusters, which is peculiar if we consider that, for example, Federated learning and Edge-computing received a very significant knowledge flow in the previous years. We argue that this decrease in knowledge flow is mostly due to the time needed for a paper to attract citations, especially outside its own cluster. Looking altogether at the three periods, notice how Cybersecurity and Edge-computing have lost influence from Blockchain over time, while Smart grids and Wireless technologies have become more reliant on Blockchain with respect to the other clusters. Moreover, we find that some clusters, such as Health, Electricity, Control theory and Governance, have received no significant influence from Blockchain in all these years, even if, Governance, for instance, has been second only to Blockchain in terms of number of papers. When looking at Fig. 6.4(b), we can see how over the years, more and more clusters have had a strong and significant impact on the future Blockchain literature. In particular, Cybersecurity, which has been one of the clusters that grew the most among all STEM clusters from the 1980s to the 2010s, has been stably the most influential cluster on Blockchain. The other top positions have instead changed from the first period considered, with Smart grids, which did not even have any influence on Blockchain at first, and Social network analysis becoming the next most important clusters. Notice also how Robot swarms and Investments have experienced an increase in knowledge flow towards Blockchain, while the opposite has happened for Telecommunication, Optimization, Governance and Health. Comparing the two plots, we find examples of only unidirectional influences between Blockchain and the other clusters. The cluster of Social network analysis, third in position since 2015 to influence Blockchain, has not been influenced by it during the same period, which is also the case of Robot swarms and Governance. A similar situation is found for Wireless technologies, that has been strongly influenced by Blockchain over time, but only in recent years has it had a small impact on it.

We have conducted a similar analysis on the Governance cluster in Fig. D.8. In this case, we consider three different periods of times: 1950–1980, that is the early stage when it was the most important cluster overall; 1981–1990, when the amount of knowledge flow




Figure 6.4: **Influences between Blockchain and the other clusters on *(de)centralization*.** (a) Change in the ranking of the most influenced clusters by Blockchain between its early period (2008-2014), its middle period (2015-2018), and its late period (2019-2020), calculated using the average knowledge flows  $K_{a \rightarrow b}(T)$ , where  $T$  is the selected period, and  $a$  is fixed to be Blockchain. (b) Change in the ranking of the clusters having most influenced the Blockchain literature (same periods as in the previous panel). In both cases, if  $K_{a \rightarrow b}(T) = 0$ , we print a circle in the corresponding gray node and use a lighter color in the respective link. Moreover, we print a star when  $0.01 < K_{a \rightarrow b}(T) \leq 0.1$ , and two stars when  $K_{a \rightarrow b}(T) > 0.1$ .

from Governance stopped to increase, still remaining among the top in terms of ranking; and 1991–2000, in which its role diminished and got surpassed by almost all other clusters by the end of the period. We do not find many noticeable differences between the first two periods. Most clusters have no significant knowledge flow from and to Governance, showing how *(de)centralization* developed independently in this cluster at first. Differently from Blockchain, the top clusters to have interactions with governance are Environment, Social network analysis and Investments. Wireless technologies, Blockchain and Robot swarms have also been influenced by Governance, but not vice-versa, apart from the sporadic case of Wireless technologies in the middle period. We can also see that the influence from Governance has increased over time on clusters like Blockchain, Optimization and Robot swarms, showing how the last years of the last century have been important milestones for the future of these clusters.

## 6.3 Discussion

In this chapter we have analysed how different topics have risen in the *(de)centralization* literature and have influenced it. By exploiting the S2AG corpus, we have shown that the literature on *(de)centralization* has exponentially increased in the past 70 years, with an author in 154 contributing to articles on the topic in 2021. Through the analysis of the evolution of knowledge flows between clusters, we have revealed that initially the different fields have had only little communication with one another to increase cross-pollinations over time, especially within STEM. Finally, we have shown how Governance has lost its leading role in favour of Blockchain, which has been the most influential cluster in the last ten years of the *(de)centralization* literature and has recently become the most productive one.

Importantly, the framework we have developed for our analysis is general and may be used to analyse the history of any concept in the academic literature. Our pipeline relies on two key methods, the multilayer hierarchical stochastic block model [279] and knowledge flows [27]. On the one hand, we employ the first one to cluster documents and words in the dataset to identify different themes and topics, using information of both citations between papers and of the words used in each document. On the other hand, knowledge flows allow us to identify significant influences between clusters over time. With the present paper, we publicly release the pipeline code to allow other researchers to perform similar analyses on other concepts.

Our study presents some limitations which also represent directions for future work. Firstly, we only consider academic papers that directly mention the word *(de)centralization* or one of its variants (e.g. “centralised”, “centralizing”, etc.). A broadened analysis could also include all articles cited by these papers, in order to further understand the roots of this topic in the different fields. Secondly, we have limited the semantic information to the document titles. Future studies could build on state of the art large scale language models and Natural Language Processing techniques to extract more information from the articles’ text (i.e., abstract and/or full text) and offer more detailed insights of their content. Moreover, one explanation behind the increase in knowledge flows between clusters could indirectly be the advent of the internet, which has made accessing research papers easier and faster than ever before. More insights on this could be gained by employing this chapter’s methodology on other scientific topics during the same time period, to see whether similar patterns can be observed. Finally, our methodology is able to identify direct flows of knowledge between two fields but misses less straightforward chains of interaction (e.g., field  $a$  influencing field  $b$ , which in turn influences field  $c$ , hence providing a possible indirect impact of  $a$  on  $c$ ). The inclusion of temporally and causally compatible higher order interactions (i.e., more than pairwise) is therefore an obvious route to improve on the current work.

Overall, our work provides new insights in the origin and evolution of the ubiquitous concept of *(de)centralization*, sheds light on the academic roots and influence of the blockchain technology, and offers a pipeline to analyse quantitatively any other concept in the academic literature. We therefore anticipate that our results will be of interest to researchers working

in a vast array of disciplines.

## 6.4 Code and data availability

All the code used for the pipeline presented in this chapter can be freely accessed and used through the Github repository available at <https://github.com/alberto-bracci/decentralization>. The data used in this work can be obtained applying the pipeline to the open-access S2AG corpus, available at <https://www.semanticscholar.org/>.

## Chapter 7

# Emergence and structure of decentralised trade networks around dark web marketplaces

Since the launch of Silk Road, the first modern dark web marketplace (DWM), in 2011 [49] millions of buyers and sellers have traded in the dark web. DWMs have become popular because their users can anonymously access them through ad-hoc browsers, such as The Onion Router (Tor) [53], and trade goods using cryptocurrencies, such as Bitcoin [3]. They offer a variety of illicit goods including drugs, firearms, credit cards dumps, and fake IDs [58]. DWMs could represent a threat for the regular economy and public health. For instance, during the COVID-19 pandemic, DWMs sold COVID-19 related goods (e.g., masks and COVID-19 tests) that were in shortage in regulated marketplaces as well as unapproved vaccines and fake treatments [16, 14, 290]. Law enforcement agencies have therefore targeted DWMs and users trading on them, performing dozens of arrests and seizing millions of US dollars worth of Bitcoin [100, 68, 291]. Despite police raids and unexpected closures, DWM trading volume has been steadily increasing and exceeded \$1.5 billion for the first time in 2020 [292].

DWM users display complex trading patterns within the marketplace environment. For example, users migrate to alternative DWMs when a DWM that they trade on closes [29, 293]. Such migration of users is aided by communication via online forums and chats on the dark web [294, 295]. However, little is known about how DWM users trade and transact *outside* the DWMs. On the one hand, some recent works have shown that a significant number of DWM users trade drugs and other illicit goods using social media platforms, such as Facebook, Telegram, and Reddit [296, 297, 298, 299, 300]. Moreover, several qualitative, interview-based studies have shown that DWM users form direct trading relationships with other users, starting user-to-user (U2U) pairs that bypass the intermediary role of DWMs [15, 301]. Past research has also found that sellers on regulated online marketplaces and social

medial platforms may decide to use intermediaries, such as Facebook groups or Instagram, to find new customers, and may start direct U2U trading with potential buyers [302]. In this chapter, we look closely at patterns of U2U trading relationships among DWM users.




Figure 7.1: **Ego and full networks.** (a) Schematic representation of an ego network surrounding a dark web marketplace (“DWM”, in red). The DWM interacts with its users (“U”, in black), which make user-to-user (U2U) pairs, represented with arrows and their respective users. (b) Multiple ego networks may be aggregated to form the full network.

The starting point for this chapter is the identification of U2U networks around DWMs. We analyse 40 DWMs for a 10-year time period spanning from June 18, 2011 to January 31, 2021. Our dataset covers all major DWMs that have ever existed, as identified by the European Monitoring Centre, Europol, the World Health Organization, and independent researchers [303, 304, 305]. Our analysis focuses on Bitcoin – the most popular cryptocurrency on DWMs [41, 42] as well as in the regulated economy [306, 307]. We focus on two kinds of transactions, occurring (i) between the user and a DWM and (ii) between two users of the same DWM. The result is 40 distinct marketplace ego networks containing user-DWM and U2U transactions, whose typical structure is depicted in Figure 7.1(a). In each network, links are directed and the arrows point at the receiver of Bitcoin. Since users often migrate from one DWM to another [29] and become users of multiple DWMs, the 40 ego networks are not isolated, and can be combined to form one full network, as shown in Figure 7.1(b).

Previous analyses of U2U trading relationships around DWMs include only two studies [15, 301] based on unstructured [15] or semi-structured [301] interviews of 17 users of Silk Road and 13 DWMs sellers, respectively. Here, we dramatically extend previous work by exploring the collective emergence and structure of U2U pairs. In section 7.2.1 we charac-

terize the U2U network emerging around DWMs. First, we observe that the U2U network, formed by all transactions between pairs of users, has a larger trading volume than DWMs themselves. We then identify stable U2U trading relationships, which represent a subset of persistent pairs in our dataset [1, 28] forming the *backbone* of the U2U network, or in other words they transact more than expected from a proper statistical null model (a detailed definition of stable pairs is presented later in the results, and in section E.3). We find that 137,667 (i.e., 1.7% out of 7.85 million total) pairs are stable, generating a total trading volume of \$1.5 billion (i.e., 5% out of \$30 billion total volume). We then explore the behaviour of users forming stable U2U pairs. We reveal that stable U2U pairs play a crucial role for marketplaces by spending significantly more time and generating far greater transaction volume with DWMs than other users. In section 7.2.3, by analysing the temporal evolution of stable pairs, we unveil that DWMs acted as meeting points for 37,192 (out of around 16 million users), whose trading volume is estimated to be \$417 million. Importantly, these newly formed pairs persist in time and transact for several months even after the closure of the DWM that spurred their formation. Finally, we observe that COVID-19 only had a temporary impact on the evolution of stable U2U pairs, which continued to increase their trading volume throughout 2020. This chapter is based on publication [V]. I contributed to this research through data preparation, methodology, study design and results interpretation.

## 7.1 Data

The dataset used in this chapter comes from the same source as the one used in chapter 5, albeit choosing a different subset of DWMs due to different constraints. The dataset contains transactions involving the 40 entities representing the 40 DWMs under consideration, which directly interact with more than 16 million other entities, who are the users of these DWMs. Users interacting with other users form U2U pairs and we include them in our dataset. The analysed dataset includes about 31 million transactions among more than 16 million users. We note that the same user can interact in multiple DWMs [29, 293]. By definition, users that interact among themselves form U2U transactions. If the pair of users interact with multiple DWMs these U2U transactions are included in all relative DWMs and counted multiple times. Therefore, the simple sum of all U2U transactions of each DWM is more than the sum of all unique U2U transactions. We count a total of 11 million transactions around all DWMs, that goes down to 9.9 million when multiple counting is avoided. Similarly, the simple sum of the single trading volumes surrounding all DWMs amounts to \$33 billion, while the overall trading volume in all unique U2U pairs is \$30 billion. Among the 40 large DWMs under consideration, 17 participated in at least one transaction in either 2020 or 2021, while the remaining 23 closed before 2020. Notably, our dataset includes Silk Road (the first modern DWM) [49], Alphabay (once the leading DWM) [308], and Hydra (currently the largest DWM in Russia) [29]. Other general statistics about our dataset can be found in Section E.2.

## 7.2 Results

### 7.2.1 Large number of U2U transactions

**Ego networks.** We start our analysis by measuring the extent of the U2U network around each DWM. The percentages of users forming U2U pairs vary across DWMs, with a median value of 38% (min 23%, max 68%). The variance in the percentage of users with U2U pairs is shown by Figure 7.2(a). It shows that the number of users with U2U pairs obeys an almost linear relationship with the number of users interacting with a DWM, having an exponent equal to 1.06 and  $R^2 = 0.969$ , see Section E.1 for details on the fitting procedure. The total trading volume users sent to the marketplace is obviously equivalent to the one they receive from it (two-sided Wilcoxon test [309]:  $W = 330$ ,  $p = 0.282$ ). Importantly, the total trading volume users sent to a DWM (and consequently the one that they receive from it) is always less than the one exchanged through U2U transactions, as shown in Figure 7.2(b).




Figure 7.2: **User-DWM and U2U transactions.** (a) Total number of users interacting with a DWM against the total number of them forming U2U transactions. The dotted line corresponds to the result of a fitted power law function. (b) Trading volume in dollars sent to a DWM compared with the total trading volume in its surrounding U2U transactions. The dashed line is the bisector and allows to easily compare the two trading volumes. (c) Total monthly trading volume sent to all DWMs and exchanged in all unique U2U pairs. We do not include the trading volume received from DWMs because it is equivalent to the volume sent to DWMs.

**Full network.** Similar results hold for the full network, confirming that the formation of U2U pairs is a pervasive phenomenon around DWMs. The total trading volume users sent to DWMs is \$3.8 billion, received from DWMs \$3.7 billion, while the volume exchanged through

U2U pairs reaches \$30 billion. In Figure E.2, we illustrate the number of transactions, trading volume, and lifespan of U2U pairs. In all cases we observe familiar fat-tailed distributions.

We then consider the temporal evolution of transactions. We look at the trading volume over time in Figure 7.2(c), where we observe that U2U transactions have consistently involved greater monthly volume than the volume sent to all DWMs since 2011. This underlines the economic importance of U2U transactions in the Bitcoin ecosystem relative to DWMs.

## 7.2.2 Behaviour of the U2U network

Henceforth, we are going to analyse users by focusing on the following groups: users who do not form stable U2U pairs; users who form stable U2U pairs, of which there are users who met outside DWMs and users who met inside DWMs (see the nomenclature in Table 7.2). We start by focusing our attention on identifying stable U2U pairs, i.e., statistically significantly persistent pairs of the U2U network. The detection of stable U2U pairs in the full network is done by using the evolving activity-driven model [1], which introduced a statistically-principled methodology to detect the network backbone against what expected from a proper null model. If a U2U pair occurs significantly more than what expected from the null model, it is labeled as stable, otherwise as non-stable. The evolving activity-driven model is an appropriate methodology for large temporal networks [28] and it is implemented in the Python 3 pip library TemporalBackbone [310], where default parameter values have been used. As input parameter, we considered the full network, comprehending transactions from/to DWMs and U2U transactions between users. For a more detailed overview of the methodology, related equations and rationale behind them, please see section E.3.




Figure 7.3: **U2U network.** The U2U network is formed by the entire set of interacting users (black and gray arrows with their respective users). Using the evolving activity-driven model [1], U2U pairs are divided in either stable (black arrows and respective users) or unstable (gray arrows and respective users).

We find 137,667 stable U2U pairs formed by 106,648 users and generating a trading volume equal to \$1.5 billion. Stable pairs produce five times more transactions per pair than non-stable pairs (two-sided Mann-Whitney-U test [311]:  $MNU = 4,58 \cdot 10^9$ ,  $p < 0.0001$ ) corresponding to a 5.34 times larger trading volume ( $MNU = 317 \cdot 10^9$ ,  $p < 0.0001$ ), see Figure E.3. Stable pairs, despite representing less than 2% of the total number of U2U pairs, generate a disproportionate amount of trading volume.




Figure 7.4: **Role of users forming stable U2U pairs.** (Main) PDFs of trading volume that users exchange with any DWMs. (Inset) PDFs of time spent by users on any DWMs. These distributions are explored for each of the 40 DWMs under consideration in Figure E.4 and E.5, respectively. Vertical lines represent median values of the respective distributions.

The high activity of users forming stable U2U pairs is not limited to the U2U network, as they are also the most active in trading with DWMs. Users in stable U2U pairs spend a median number of 41 days on DWMs versus a median of only one day for users without stable pairs. The two resulting distributions are significantly different (two-sided Kolmogorov-Smirnov test [312]:  $KS = 0.673$ ,  $p < 0.0001$ ), see the inset of Figure 7.4. When we look at the trading volume with DWMs, we find qualitatively similar results. Users in stable U2U pairs transact a median of \$400 with DWMs, while other users transact only \$56. The two resulting distributions are significantly different ( $KS = 0.438$ ,  $p < 0.0001$ ), see Figure 7.4. These results hold not only for full network but for every DWM in our data, see Figure E.4 and E.5.

### 7.2.3 U2U network evolution

**Formation of U2U stable pairs.** Having mapped the behaviour of stable pairs, we now consider their temporal evolution. More specifically, we ask: How do stable pairs form? Do DWMs spur their creation? One possible hypothesis is that users meet for the first time while active on a DWM, i.e., after they have both traded with that DWM. This can be considered as a plausible, and conservative, proxy for users who met inside a DWM. We determine whether U2U pairs meet while active on a DWM by looking at the time occurrence of their first U2U transaction. This transaction can occur at three different moment in time. (i) At  $t = t_1$ , before both users interact with the same DWM (occurring at  $t = t_2 > t_1$  and  $t = t_3 > t_1$ , respectively), as shown on the left hand side of Table 7.1. (ii) At  $t = t_2$ , when only one user has interacted with a specific DWM and the other user will do so at a later time, as in the middle column of Table 7.1. (iii) At  $t = t_3$ , when both users have interacted with the same DWM, as in the right column of Table 7.1. We classify these three chain of events in two groups. One group includes all pairs that meet outside any DWMs, which includes case (i) and case (ii), and the other group users that meet inside a DWM, described by case (iii). This last case constitute a conservative proxy for users that meet who met inside a DWM. The proxy admits the possibility of false positives, since it consider users who met inside the same DWM without having interacted on it, as well as false negatives, since it does not take into account users who met inside a DWM without having ever interacted on it. The latter is arguably more significant, since it is possible that only one of the two users (the seller) has actually engaged in transactions with the DWM, while the other user, after seeing the seller’s profile on a DWM, has established a direct contact. To recap, the definitions of all considered groups can also be seen in table 7.2




Table 7.1: **Formation mechanism of stable U2U pairs.** We compare the time in which the first transaction between a pair of users occur with the time in which these users interact with the same DWM. Each row in the figure indicates a possible temporal sequence, which we classify in two groups: users who met outside the DWM (first two columns) and users who met inside the DWM (last column).

Group	Description	Number of users
1.	Users who do not form stable U2U pairs	15,871,206
2.	Users who form stable U2U pairs	106,648
2a.	Users who met outside DWMs	88,828
2b.	Users who met inside a DWM	37,129

Table 7.2: **Nomenclature.** Definitions of all groups the users are divided to based on their behaviour. Number of users in each group is given in the last column.

A total of 37,129 users have met at least one other user inside a DWM. Their trading volume is about \$417 million, and the percentage of users who met inside a DWM is proportional to the trading volume sent to DWMs (Spearman [313]:  $C = 0.805$ ,  $p < 0.0001$ ), see Fig E.6, meaning that large DWMs are more likely to favour the encounter of users than smaller DWMs. Importantly, users who met inside a DWM transact more than those meeting outside them. In particular, users who met inside a DWM trade a median of \$2,212 between themselves, almost twice the \$1,379 for users meeting outside the DWM ( $MNU = 1.863 \cdot 10^9$ ,  $p < 0.0001$ ). Moreover, users who met inside a DWM tend to transact for significantly longer with median of 61 days than users meeting outside with a median of 50 days ( $MNU = 2.099 \cdot 10^9$ ,  $p < 0.0001$ ).




Figure 7.5: **Resilience of stable U2U pairs after DWMs closure.** Trading volume of U2U pairs surrounding active DWMs. (Main) U2U pairs meet who met inside aa DWM. (Inset) U2U pairs meet outside them. Curves indicate the median value while bands represent the 95% confidence interval. Day zero corresponds to the day when the market closed. Negative and positive numbers indicate the days prior and after the closure, respectively. Only the 33 DWMs closed are considered in the analysis.

**Resilience of U2U stable pairs.** Thus far, we have shown that users involved in stable trading relationships are also very active on DWMs, where they may meet new trading partners. But are DWMs and the U2U network truly interdependent? In particular, do stable pairs need the DWMs to survive? To answer these questions, we look at market closures, previously investigated to show how active users migrate to other existing DWMs [29]. Our dataset includes 33 closure events, which we study independently from one another by considering the evolution of the respective 33 marketplace ego networks. We find that unstable U2U pairs sharply stop interacting immediately after the DWM closure, and therefore their existence is highly sensitive to the presence of the DWM. On the other hand, the trading volume of stable U2U pairs is only marginally affected by the disappearance of the DWM. As a result, while prior to DWM closures unstable U2U pairs generate an overall trading volume that is 10 times higher than that of stable U2U pairs (since unstable pairs are far more prevalent), within a few weeks after DWM closures the pattern is reversed: stable U2U pairs generate more trade volume than unstable U2U pairs. Indeed, trading patterns of stable pairs are not significantly influenced by the sudden DWMs closure, and they very slowly decay over time, see Figure 7.5.




Figure 7.6: **Temporal evolution of stable pairs.** (a) Monthly number of new stable U2U pairs created. (b) Monthly trading volume of stable U2U pairs.

We have shown that the U2U network is resilient to short-lasting external shocks, namely the closure of a marketplace, and it does not need the centralised structure of DWMs to survive. What about long-lasting systemic stress? To answer this question, we consider the impact that the COVID-19 pandemic has had on the evolution of stable U2U pairs. Previous studies reported that COVID-19 had a strong impact on DWMs, with reported delays and damage to the shipping infrastructure due to border closures [314, 315]. We start by investigating the number of new stable U2U pairs and their trading volume. Users in stable pairs meeting both inside and outside DWMs have been growing over the last two years, since the shutdown of AlphaBay [68], the largest DWM at the time. In 2020, a total of 6,778 pairs of users in stable pairs met inside a DWM, more than double the amount of 2018 and 2019 respectively, see Figure 7.6(a). Pairs of users in stable pairs meeting inside a DWM traded for a total of \$145 million in 2020, which corresponds to almost six times the amount of 2018, see Figure 7.6(b). We see similar trends for stable U2U pairs meeting outside any DWMs. The impact of the COVID-19 pandemic has, however, had different phases, determined by the number and level of measures introduced around the world. For users in stable pairs who met both inside and outside DWMs, we find that during the first lockdowns in 2020 trading volume fell with respect to January of the same year, suggesting that they were negatively impacted by COVID-19 restrictions. After that, trading volume sharply increased over all 2020, see Figure E.7. The number of stable U2U pairs created each day was, however, steady over time during 2020, even though more U2U pairs were

created compared to the same period of 2019, see Figure E.8. Overall, stable U2U pairs have shown resilience to the systemic stress caused by COVID-19, suggesting, once again, that these trading relationships are fundamentally independent from the underlying DWMs.

### 7.3 Discussion and Conclusion

In this chapter, we revealed the prevalence and structure of a large network of direct transactions between users who trade on the same DWM. We showed that some of the links of this user-to-user (U2U) network are ephemeral while other persist in time. We highlighted that a significant fraction of stable U2U pairs formed as their members were trading with the same DWM, suggesting that DWMs may play a role in promoting the formation of stable U2U pairs. We showed that the relationships between users forming stable pairs persist even after the DWM shuts down and are not significantly affected by COVID-19, suggesting overall resilience of stable pairs to external shocks.

Our study has several limitations. In particular, our dataset does not include any attributes related to either users or their Bitcoin transactions, such as, whether the transaction represents an actual purchase or not. Moreover, we do not have information about which users trade with other users on the same DWM. Finally, our coverage of DWMs, albeit extensive, may lack information on other DWMs where users could have met.

Our work has several policy implications. Our findings suggest that DWMs are much more than mere marketplaces [316]. DWMs are also communication platforms, where users can meet and chat with other users either directly – using Whatsapp, phone, or email – or through specialised forums. These direct interactions may favour the emergence of decentralised trade networks that bypass the intermediary role of the marketplace, similar to what is currently happening on Facebook, Telegram, and Reddit [296, 302, 297, 298, 299, 300], where users post products, negotiate item prices, and then trade directly without an intermediary. We estimate that the trading volume of U2U pairs meeting on DWMs is increasing, reaching a peak in 2020 (during the COVID-19 pandemic). By contrast, trading volume on DWMs was negatively affected by COVID-19, mainly due to shipping delays [314, 315]. The reasons for the differential impact of COVID-19 on U2U trading vs. DWM trading are difficult to pin down. One hypothesis is that U2U pairs managed to find better shipping logistics; another hypothesis is that they were seen as a safer way to trade than DWMs at a time of crisis.

Our results also support recent recommendations of paying attention to single sellers rather than entire DWMs [317]. Law enforcement agencies, however, have only recently started targeting single sellers. The first operation took place in 2018 and successfully led to the arrest of 35 sellers [318], while the largest operation to date occurred in 2020 and led to 179 arrests in six different countries [319]. Our study indicates that a much higher number of highly active DWM users, to the order of tens of thousands, is involved in transactions with other DWM users. Moreover, our analysis paves the way for a deeper understanding

of U2U transactions in online marketplaces. Recent results have shown that transaction networks and activity on DWMs and regulated online marketplaces share several robust macroscopic properties [320]. One might therefore hypothesise that U2U trading is also a prevalent feature of regulated online marketplaces. While data on U2U transactions is far harder to obtain (as these transactions might involve a variety of commercial methods), there is clearly a need to better understand the dynamics and structure of trading relationships beyond what is observable on a specific online marketplace.

Overall, our study provides a first step towards the understanding of how users of DWMs collectively behave outside organised marketplace. We believe that the results might suggest to researchers, practitioners, and law enforcement agencies that a shift in the attention from the evolution of DWMs to the behaviour of their users might facilitate the design of more appropriate strategies to counteract the online trading of illicit goods.

## Chapter 8

# Heterogeneous rarity patterns drive price dynamics in NFT collections

Throughout 2021, the NFT market grew by more than 61,000%, starting from a monthly sale volume of 8,072,866 USD in January 2021 to 4,968,834,938 USD in January 2022 [321]. NFT was Collins Dictionary’s word of the year for 2021 [322]. NFT collections are groups of NFTs that share common features, such as visual aspects or the code that generated them [323]. They have been a driving force for the booming NFT market [17, 324]. In the prominent case of generative art, NFTs are associated to (virtual) objects made using a predetermined system, typically an algorithm, that often includes an element of chance [325]. To be concrete, CryptoPunks is a collection of 10,000 unique images of pixelated human faces algorithmically generated [326], while Bored Ape Yacht Club contains 10,000 profile pictures of cartoon apes that are generated by an algorithm [327]. Their market capitalization is 834M USD and 1.2B USD as of June 2022, respectively [328].

NFTs in a collection are most often distinguishable from one another. For example, CryptoPunks have a gender (6,039 male and 3,840 female) and – as for many other collections – a number of *traits* that distinguish them. So a punk can have, or not have, a “Top Hat”, a “Red Mohawk”, a “Silver Chain”, or “Wild White Hair” among other possibilities. Furthermore, while most CryptoPunks are humans, there are also 88 Zombies, 24 Apes, and 9 Aliens in the collection. CryptoPunks are not equivalent according to the market. The most expensive CryptoPunk to date was sold for 23.7 million USD on February 12, 2022 [329], despite the average price of a punk being “only” 138,179 USD (see also [330]). A similar picture holds for Bored Apes, with the most expensive one traded for 3.4 million USD on October 26, 2021 [331], vs an average price of 48,836 USD.

An hypothesis to rationalise these differences in price considers rarity. The heterogeneous distribution of traits among NFTs make some of them more rare than others (see Figure 8.1), and scarcity is attractive for collectors [332, 333, 334, 335]. However, despite some evidence that rarity and aesthetic preferences do play a role in the case of CryptoPunks [336, 330], a

comprehensive analysis of the role of rarity on the market of NFTs is still lacking.




Figure 8.1: **Illustration of rarity in a collection.** Demonstration of the visual difference between rare and common NFTs using the example of CryptoPunks. CryptoPunk #2547 (on the right) is the least rare, as it has traits that appear frequently in the collection (i.e., the bandanna and the earring). CryptoPunk #8348 is the rarest in the collection, mostly since it is the only one with seven non-null attributes. Rarity scores are not normalised. After normalisation, the total rarity score for punk #2547 is zero, while the one for #8348 is 100 (min and max of the collection, respectively). In the bottom left corner, we show, as an example, the rarity score of traits associated to the “Facial Hair” attribute.

In this chapter, we carry out a systematic investigation of how the rarity of NFTs impacts their market behaviour. We focus on within-collection rarity using the definition proposed in the platform rarity.tools [328]. Our dataset describes the rarity of 410 collections listed on OpenSea, containing a total of 1,479,020 NFTs that were exchanged 3,775,040 times between January 23, 2018 and June 6, 2022. In section 8.2.1 we show how the NFT collectibles markets has grown in the recent past, characterizing the impact of different collections in this phenomenon. In section 8.2.2, we characterise trait distributions and investigate how they impact NFT rarity. Then, in section 8.2.3 we analyse transaction data and find that rarity positively correlates with NFT prices and negatively correlates with number of sales. Finally,

in section 8.2.4, we find that rarity also positively correlates with return on investments (ROIs), while negatively correlates with risk, quantified as the likelihood of a negative ROI. The breadth of our analysis suggests this market behaviour is likely to be genuinely self-organised. At the same time, our results could inform further research aimed at establishing how to optimally design collections, as well as effective trading strategies for the NFT market. This chapter is based on publication [V]. I contributed to this research through methodology, study design, data analysis and result interpretation.

## 8.1 Background, data and methods

### 8.1.1 Glossary of key terms

*Attributes and traits.* Attributes refer to human-readable characteristics of an NFT. In generative art, for example, they usually relate to visual properties of items. Attributes can take one among several values. For example, in the CryptoPunks collection, every item has the attribute “type” that can take one among the following traits: “Male”, “Female”, “Zombie”, “Ape” or “Alien”. CryptoPunks have also attributes that capture the presence of any accessory, such as earrings or bandanas. For the remainder of this study, we refer to the value taken by an attribute as the *trait*.

*Collections.* A collection is a group of NFTs whose associated digital items share common features. When minting an NFT, a creator can include the corresponding item within a collection. In generative art, for example, items of a collection are created by the same generative algorithm.

*Marketplaces.* Creators and collectors meet in online marketplaces to trade NFTs. The largest of these markets, OpenSea [337], enables any creator to sell their NFTs and, at the moment of writing, it offers 44 million NFTs [338]. Other marketplaces feature a curated selection of creators (e.g., Foundation [339], SuperRare [340], Nifty Gateway [341], Feral File [342]). NFTs are auctioned on these marketplaces, where the NFT can be sold to the highest bidder or with a declining price, depending on the kind of auction. After an NFT is minted on a marketplace – i.e., it is converted into a digital asset on the blockchain – it can be put up for auction. Typically, the first transaction, from the creator of the NFT to the first user, is different in nature from the subsequent trades (e.g., the first user is often not able to select a specific NFT from a collection [343]).

### 8.1.2 Dataset

Our dataset includes 3,775,040 sales, taking place on the Ethereum blockchain, of 1,479,020 NFTs from 410 collections, including 61 of the top 100 collections by sales volume according to CoinMarketCap [344]. The list of collections considered in this study can be found in section F.3. The dataset was built by considering all collections we could automatically match (by name) between rarity.tools [328] - a website dedicated to ranking collectible NFTs,




Figure 8.2: **Characteristics of collections.** The probability distribution of a) the total number of NFTs within the collection; b) the total number of attributes in the collection; c) the total number of traits in the collection.

also sometimes called Profile Picture NFT projects (PPF), by rarity - and the Opensea market. From the latter, we collected the release date, NFT traits and all sales concerning these collections that took place between January 23, 2018 and June 6, 2022. To avoid spurious effects, we only considered user-to-user transactions, where buyer and sellers are both aware of the precise identity of the traded NFT (i.e., we discarded the initial creator-to-user transactions). In the following, we refer to the first user-to-user transaction as “primary” sale, and to all subsequent transactions as “secondary” sales. Where not specified, by “sales” we consider both primary and secondary sales.

Collections in our dataset have on average 7,554 NFTs. There is, however, wide heterogeneity across collections: the standard deviation of the distribution is 194.64, and the mode is 10,000 NFTs (see Figure 8.2a). The number of attributes in a collection is  $11.1 \pm 0.91$ , where the reported error corresponds to the standard deviation of the distribution (see Figure 8.2b). As for the number of traits, the average is equal to 415.1, with a standard deviation of 97.6 (see Figure 8.2c). On average, an attribute within a collection has 37.4 different traits. More information about the algorithms used to assign traits to an NFT can be found in section F.2.

### 8.1.3 Rarity

The rarity of a trait is quantified as the fraction of NFTs within a collection having this trait. This value is indicated on OpenSea’s sale page. For a collection containing  $N$  NFTs, the *trait rarity score*,  $R_t$ , for a trait  $t$  shared by  $r$  NFTs is defined as:

$$R_t = \left(\frac{r}{N}\right)^{-1} \quad (8.1)$$

To quantify the overall rarity of an NFT within a collection, we consider each trait independently and define the *NFT rarity score*,  $R_{NFT}$ , as the sum of the rarity scores of each one of its traits, that is

$$R_{NFT} = \sum_t R_t. \quad (8.2)$$

In order to compare this score between collections, we then normalize the scores within a collection with a min-max normalization. For a collection with a maximum and a minimum rarity score  $R_{max}$  and  $R_{min}$  respectively, the normalised rarity score  $R_{norm}$  is given by  $R_{norm} = 100(R - R_{min}) / (R_{max} - R_{min})$ . By doing so, every NFT ends up with a normalised rarity score between 0 (least rare) and 100 (rarest). All the analyses presented in the main text of this article are based on the NFT rarity score.

Finally, we also consider *the NFT rarity rank*, where the rarity rank of an NFT is given by the trait rarity rank of its rarest trait. This metric will allow us to quantify the effect of a rare trait on the market behaviour of an NFT, regardless of its other traits. Analyses based on the NFT rarity rank can be found in section F.5.1.

Further information on measuring NFT rarity, including a detailed discussion of the above measures, can be found in [345].

## 8.2 Results

### 8.2.1 Market Growth




Figure 8.3: **The collectible market boom.** 30-day rolling average of the daily number of primary sales (blue line), secondary sales (orange line) and new buyers (i.e., new wallet addresses buying an NFT for the first time, green line). Inset: 30-day rolling average of the daily trade volume (in USD).

We start by investigating the evolution of NFT sales in our dataset over time. We find that the interest in the collections remained stable until the end of 2020, then started to gain

traction in 2021, especially in terms of available NFTs on the market (see Fig. F.1). The number of primary sales grew from an average of 14 daily sales in January 2021 to 784 sales every day in March 2022, when the market peaked, implying a percentage growth of 5,500% (see Figure 8.3). Similarly, secondary sales grew by 110,177%, starting from 9 sales/day in January 2021 and reaching 9925 sales/day in March 2022. Interestingly, around October 2021, the number of secondary sales started to exceed the number of primary sales, a trend that still holds at the moment of writing. This surge in activity led to a growth of daily volume of trades of 18,520% between January 2021 and March 2022 (see Figure 8.3 inset), and attracted new users. The number of new buyers increased by 41,755% in 2021. These results indicate an overall growth of the popularity of NFT collections on OpenSea, both with respect to the size of the NFT community, and to the total market value.

Different collections contributed to varying extents to the growth of the collectible NFT market. Figure 8.4 shows the distribution of key market properties across NFT collections: total number of sales per collection (Figure 8.4a), total traded volume per collection (Figure 8.4b) and collection items median sale price (Figure 8.4c).

Collections are widely heterogeneous with respect to market properties. 25.6% of the collections have generated less than 1,000 sales, whereas 17.1% have generated more than 10,000 (see Figure 8.4a). Further, 43.9% of the collections had a total trade volume below a million dollars, whereas 3.64% generated more than a hundred million dollars of sales on the marketplace (see Figure 8.4b). The success of a collection can also be measured by looking at the median price at which its NFTs are sold on OpenSea. For 18.3% of the collections, the median sale price is lower or equal to a hundred dollars, whereas it is higher than a thousand dollars for 12.9% of the considered collections (see Figure 8.4c). These findings indicate that collectibles NFT do not meet the same success on OpenSea, a claim that is supported by the infamous success stories of a few collections, whereas the others quickly become a flop on the platform [346].




Figure 8.4: **Cumulative distribution of key market properties across collections.** a) Total number of sales per collection. b) Total trade volume per collection. c) Median sale price per collection. In inset: Distribution of each market property.

## 8.2.2 Quantifying rarity

We quantify the distribution of rarity scores for items within the same collection. As an example, Figure 8.5 shows the distribution of rarity for three popular collections, namely CryptoPunks, Bored Ape Yacht Club, and World of Women.

For CryptoPunks, the median rarity score is 0.82, with only one of the 10,000 CryptoPunks having a rarity score above 75, whereas 99.7% of the tokens have a rarity score below 10 (see Figure 8.5a). Moreover, as most of the CryptoPunks have a low rarity score, the least rare ones are aggregated into two bins, whereas the rare one occupies the only bin with a high rarity score within the collection. The median rarity score for Bored Ape Yacht Club is 20.3, and 26 apes (i.e., 0.26% of the collection) have a rarity score above 75. The distribution is skewed towards lower rarity scores, with 68.2% of the assets with a rarity score below 25, among which 8.23% fall below a rarity score of 10 (see Figure 8.5b). The profile for the World of Women collection is also not as heterogeneous as that of CryptoPunks; it has a median rarity score of 14.8 and only 24 assets (0.24% of the collection) have a rarity score above 75. 87.3% of the tokens have a rarity score below 25, and 19.9% of those lie below a rarity score of 10 (see Figure 8.5c). To generalize these observations, we calculated the Spearman rank correlation coefficient between the rarity bin and the number of NFTs by rarity bin. A negative value of the correlation coefficient indicates that the higher the rarity score, the lower the supply of NFTs is within the considered collection. Like the three example collections in Figures 8.5a-c, 96% of the collections in our dataset have a Spearman rank  $r \leq 0$ , as shown in Figure 8.5d, where the violin plot represents the probability distribution of the Spearman rank correlation by collection. We compare the ability of 6 different statistical distributions, namely the exponential, power-law, uniform, cauchy, log-normal and levy distributions, to capture the distribution of rarity for each collection, using the Akaike model selection method [2] (see section F.4 for more details). We find that, among the distributions considered, 90% of the collections are best described by a log-normal distribution (with  $\langle \mu \rangle = 0.91 \pm 0.16$ , see Fig.F.2), only 7% by an exponential, 1% by a uniform function and the rest by heterogeneous distributions such as power-laws or Levy (for a visualization of a sample of these distributions, see Fig. F.3).

The same correlation analysis performed using the rarity rank confirms our results (see Fig. F.4) In the following, we will focus on NFTs rarity score, because this metric takes into account all the traits associated with an NFT, and is therefore more suitable to quantify NFTs properties and rarity. All the following results are replicated using trait rarity rank as a robustness check (see section F.5.1)

Our analysis indicates that the distribution of the rarity within a collection is heterogeneous, thus leading to a situation where rare NFTs are genuinely scarce on the marketplace. Notice that while this may seem trivial (“rare items are fewer than common items”), the distribution of trait rarity, and in turn their combination in single NFTs could in principle generate a wide range of distributions of NFT rarity, including homogeneous ones. The heterogeneously rare traits could indeed be distributed among NFTs in such a way that the

NFTs rarity is more homogeneous, by assigning the trait values in such a way that eq. 8.2 gives similar scores to different NFTs with different traits. One must also remember that traits are what is algorithmically generated by the collection’s creators, while the NFT rarity is instead an emergent behavior of the system.




Figure 8.5: **Rare NFTs are scarce within a collection.** a-c) Distribution of the rarity score of the NFTs within three collections: CryptoPunks (a), Bored Ape Yacht Club (b), and World of Women (c). d) Violin plot of the Spearman Rank correlation computed between the rarity score and the number of NFTs with that score. 96% of the collections have a Spearman rank  $r \leq 0$  (black dashed lines).

### 8.2.3 Rarity and market performance

To measure the relationship between rarity and market performance, we compute the rarity score of each NFT, and we split the assets into quantiles with respect to their rarity score to analyse collections individually. We then compare the median sale price across quantiles. We are using quantiles to ensure that NFTs within a collection will be evenly balanced between each bin, as to avoid having a collection skewing the results in the aggregated analysis, by having all of its NFTs concentrated in a single bin. For the individual collections analysis, NFTs are partitioned into twenty quantiles, whereas 100 quantiles are used when aggregating the collections together.

First, we consider the relation between market behaviour and rarity for three exemplar collections, CryptoPunks, Bored Ape Yacht Club, and World of Women (see Figure 8.6a-c). We observe that the median sale price at which NFTs are auctioned is relatively constant for

the most common NFTs in each collection (rarity quantile smaller than 10), and then increase sharply for the rarest NFTs (rarity quantile larger than 10, see Figure 8.6a-c). These findings are robust, and are observed also when we consider NFTs in all collections (see Figure 8.6d). We notice that the median sale price is relatively flat for the 50% least rare NFTs, before increasing by 195% for the 10% rarest NFTs. More strikingly, the median sale price for the 90% least rare NFTs is equal to  $298 \pm 3.2$  USD, and rises to 1,254 USD for the 1% rarest NFTs. Focusing on the top 10% rarest NFTs, the relationship between the median sale price  $p$  and the quantity  $(100 - q)$ , where  $q$  is the rarity quantile, is well described by a power law  $p \sim (100 - q)^\alpha$  with exponent  $\alpha = -0.55$  (see Figure 8.6 inset). This result indicates a strong relationship between NFT rarity and median sale price.

On the other side, we find that rare NFTs are not sold as frequently as common ones on the marketplaces. By looking at the individual collections, we see that the average number of sales decreases as we increase the rarity of the NFTs we are considering (see Figure 8.6e-g). Regarding the average number of sales, by aggregating all collections together, we find that the number of sales decreases for rarer NFTs. In particular, the 1% least rare NFTs are sold, on average, 10.8% more than the 1% rarest ones (see Figure 8.6h).

In order to check that this behaviour holds when considering a shorter time span within OpenSea's lifetime, we performed the same analysis by considering only sales happening during the third quarter of 2021 (see Fig. F.9) and the fourth quarter as well (see Fig. F.11). Our findings are also robust by considering the sale price in ETH rather than in USD (see Fig. F.7), and by discarding the rarest and least rare NFTs from each collection (see Fig. F.13). Moreover, we notice a similar pattern when quantifying the rarity of the NFTs with the NFT rarity rank instead of the NFT rarity score (see Fig. F.5).




Figure 8.6: **Rare NFTs have a higher financial value and circulate less on the marketplace.** Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity  $(100-q)$ , where  $q$  is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections.

## 8.2.4 Rarity and return on investment

NFTs can be purchased and later put on sale again on the marketplace. An NFT owner is free to set an initial price to an auction, and to transfer the ownership of the NFT to the highest bidder. As such, NFTs which have been minted years ago, such as the CryptoPunks, can still be purchased on OpenSea. The results shown in Figure 8.6 indicate that, within a collection, the rarest NFTs are typically sold at a higher absolute price than the least rare ones on the market. However, this fact does not necessarily imply that the return on investment of secondary sales is positive, as it does not take into account the price at which the asset was initially purchased before being auctioned again. To study whether the correlation between rarity and price strengthens as a token keeps being exchanged on the market, we computed the return  $v$  of the  $k^{\text{th}}$  sale of an NFT as:

$$v = \frac{P(k) - P(k-1)}{P(k-1)}, \quad (8.3)$$

where  $P(k)$  is the price that was paid for the NFT for its  $k^{\text{th}}$  sale. A positive return indicates that the NFT was sold at a higher price than the one it was bought for, whereas a negative return represents a financial loss for the seller.

Figure 8.7a shows the median return computed when aggregating all collections by rarity quantile. We find that the rarest NFTs have a much higher median return, whereas the value is almost constant in the first half of the curve. Focusing on the top 10% rarest NFTs, we observe that the relationship between the quantity  $(100 - q)$ , where  $q$  is the rarity quantile, and the median return  $v$  is well described by a power law  $v \sim (100 - q)^\alpha$ , with an exponent  $\alpha = -0.29$  (see Figure 8.7 inset). The median return is relatively flat around  $0.24 \pm 0.001$  for the 50% least rare NFTs, thus indicating no noticeable advantage for an NFT to be one of the least rare assets of the collection or an average one in terms of rarity, whereas the median return grows by 105% within the top 10% rarest NFTs. Finally, we study the relation between NFT rarity and the probability to generate negative returns. We observe that, on average, rarer NFTs are less likely to generate negative returns (see Figure 8.7b). The fraction of sales generating negative returns is equal to  $34.6 \pm 0.58\%$  for the 50% least rare NFTs, but drops from 30.5% to 22.9% within the top 10% rarest NFTs, i.e., a decrease of 24.9%. These results also hold by only considering the sales happening during a shorter time period, such as the third quarter of 2021 (see Fig. F.10) and the fourth quarter (see Fig. F.12). The same analysis has been performed by considering the sale prices in ETH (see Fig. F.8) and by discarding the rarest and least rare NFTs of every collection (see Fig. F.14). These results are also robust when using the NFT rarity rank to measure the rarity of an NFT rather than the rarity score (see Fig.F.6).




Figure 8.7: **High rarity leads to higher returns, and a lower chance of a negative return.** a) Median return in USD by rarity quantile. Inset: median return against the quantity  $(100 - q)$ , where  $q$  is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile.

### 8.3 Discussion

We have quantified rarity in 410 NFT collections and analysed its effect on market performance. Rarity is a fundamental feature of NFTs belonging to a collection because (i) it allows users to categorise NFTs on the traditionally market-relevant axis of scarcity and (ii) it is based on human-readable, easy to identify, traits that creators assign to NFTs. We have found that the distribution of rarity is heterogeneous throughout the vast majority

of collections. We have shown that rarity is positively correlated with the sale price and negatively correlated with the number of sales of an NFT, with the effect being stronger for the top 10% rare NFTs. Last, we have shown how rarity is associated with higher return of investment and lower probability of yielding negative returns in secondary sales.

The finding that most rarity distributions are heavily heterogeneous, with few very rare NFTs, is interesting since in principle more homogeneous distributions would be possible. The ubiquitous nature of this pattern may indicate either that creators deliberately choose heterogeneous distributions (design perspective) or that heterogeneous distributions help make a collection successful and therefore are dominant in our sample of actively traded distributions (evolutionary perspective). While information on the rationale behind rarity distributions is hard to retrieve [347], the design and evolutionary explanations could have fuelled one another over time, with creators embedding rarity out of imitation of successful pre-existent collections. In this perspective, our results could help to further improve the design of NFT collections.

From the point of view of trading, it is important to highlight that our results concern genuinely emerging properties of the NFT market, since we only considered user-to-user sales. In doing so, we discarded the very first creator-to-user sales, which are often based on lotteries that prevent users to select what NFT to buy [343]. We found that while the impact of rarity is particularly strong for – and among – the rarest NFTs, which are thus genuinely non-fungible according to the market, its influence propagates to a large number of somehow rare NFTs (see Figure 8.6g, inset and Figure 8.7a, inset). Most common NFTs, on the other hand, appear to behave more uniformly in the market, which appears to consider them essentially “fungible”. Overall, we anticipate that our results in this context may help inform the decisions of users interested in the financial aspects of NFTs.

Our study has limitations that future work could address. First, our dataset is limited to collections available on OpenSea, the biggest NFT market, and sold on the Ethereum blockchain. A natural extension would cover other platforms (potentially on other blockchains) and different types of NFTs, assessing whether rarity has the same effects on other kinds of NFTs such as those related to gaming and the metaverse. Second, we used the rarity score to quantify the rarity of an NFT. While this method does take into account every trait associated with an NFT, it does not consider possible combined effects stemming from the combination of multiple traits (e.g., two common traits for a collection might be present together in just one NFT, making it very rare). Future work could assess whether such second-order effects do play a role on the market performance of NFTs. Third, we considered traits as they are encoded in the NFT metadata and reported on rarity.tools, limiting the analysis to collections where such metadata are available and consistently recorded. Future work making use of computer vision techniques to extract human readable attributes from visual information of NFTs would yield to larger datasets and assess whether also less “official” visual traits, potentially shared by NFTs in multiple collections and where previously developed metrics might help [348, 349], might play a role on the NFT market. Finally, while this work has focused on how rarity affects NFT market success, a natural extension

of the work should focus on how buyers behave with respect to rarity.

# Chapter 9

## Conclusions

This thesis was inspired by recent advances in blockchain-based ecosystems and by the availability of unique comprehensive large scale datasets giving us the possibility to improve our understanding of these systems and collective human behaviour in general. Compared to previous research, our results can be summarised along four main points encompassing the main chapters of this thesis. Firstly, we have improved our understanding of multiple blockchain based ecosystems, including Dark Web Marketplaces and online marketplaces, the user to user trading network around DWMs, the NFT market and the scientific literature on decentralization. Secondly, our analysis has covered the largest possible period of time, including the whole history of DWMs (and user to user interactions around them), the NFT market boom of 2018-2022 and the blockchain and related literature since the 1950s. Third, we have employed a complex systems approach in our research, improving the understanding of the interplay between these systems, society, technology and public attention using state of the art method from applied mathematics and statistical physics, also bringing forward methodological novelties in the form of network formation models and data analysis pipelines. Finally, we have analysed novel unique large scale datasets, both obtained by private companies (when impossible to openly gather them on our own) and collected from openly available sources, which show the potential of using blockchain related data to study the behaviour of complex decentralized socio-technical systems and human behaviour in general.

Our research addressed five main questions which had attracted the interest of the scientific community. Our main findings can be summarised in five main points.

*Dark Web Marketplaces quickly reacted, and then further adapted, to the external shock caused by the COVID-19 pandemic, swiftly offering, and adapting their offer in time, goods in shortage in the regulated economy or goods under strong public attention, like chloroquine, face masks and COVID-19 vaccines.*

In chapter 3 we have analysed the offer on DWMs between January and November, 2020.

We have shown how the markets swiftly reacted to the pandemic by quickly starting to offer products in shortage on the regulated economy such as PPE, and other COVID-19 related products like guides on scamming and COVID-19 test records. We have then linked the offer on DWMs with public attention, showing how the number of listings of PPE and hydroxichloroquine correlated with public attention on these products, measured using data from Twitter and Wikipedia. Finally, we have shown how the pandemic had an indirect impact on DWMs, as measured by mentions of sales, lockdowns and delays correlating with major COVID-19 pandemic events.

In chapter 4, we have extended the analysis by studying how DWMs adapted in the following period of the pandemic, from November 2020 to July 2021, and by uncovering the impact on non COVID-19 related listings. First, we have shown how DWMs have kept adapting to the changing situation of the pandemic, offering COVID-19 vaccines first, and then fabricated proofs of vaccinations (sometimes known as COVID-19 passports) later. Second, we have shown how these listings have actually replaced previously offered COVID-19 listings, such as PPE, which were now easily available on the regulated economy. Finally, we have shown how drugs have been the most affected non-COVID-19 related category of goods. Overall, by analysing the largest available dataset of DWM listings to the date of publication of this work, we have shown how DWMs form a reactive blockchain-based ecosystem able to swiftly and quickly react and adapt to an external shock such as the COVID-19 pandemic, while also adapting to public attention on specific goods.

*Dark Web Marketplaces and regulated online marketplaces share macroscopic buyer-seller network properties, captured by a human behaviour network formation model based on preferential attachment and buyer memory.*

In chapter 5 we have analysed a unique comprehensive dataset of transactions happening on DWMs and regulated online marketplaces between 2010 and 2021, available thanks to partnership with private companies. We have shown how both kinds of platform share the same common macroscopic buyer-seller network properties, not only between different platforms of the same kind (e.g. different DWMs), but also between regulated and unregulated markets, a surprising result considering the many differences between the blockchain-based markets and the regulated one. We have then investigated the microscopic dynamics of buyers, showing how their behaviour is affected by the memory of past interactions, modelled in the same way as different social interactions. Finally, drawing on the empirical observations, we have proposed a complex network formation model reproducing the main properties of the buyer-seller networks, showing the prominent role of memory and preferential attachment mechanisms in shaping the system. Overall, our results point to the role of human behaviour in driving the properties of online marketplaces, and blockchain-based Dark Web Marketplaces in particular, hinting at the fact that the specifics of these systems do not contribute to shaping their macroscopic collective behaviour.

*The concept of Decentralization has had a decentralized evolution, with Governance being the most influential field in the past, replaced by Blockchain since its inception.*

In chapter 6 we have studied the origin and evolution of the concept of Decentralization, pillar of blockchain technologies, through the lenses of the academic literature. First, we have shown the rising interest on the topic, with 1 author in 154 contributing to the literature in 2021. Then, we have mapped the academic literature on the topic by using a multilayer stochastic blockmodel, incorporating semantic and citation information, showing how different fields have contributed to the literature. By using the methodology of knowledge flows, we have been able to show how Decentralization has had a decentralized evolution itself, with little exchange between different fields in the origin years. Finally, we have shown how Governance and Blockchain are the two paradigmatic fields in the topic, the former being the most influential until the 1990s, the latter being the most influential in recent years. Overall, our results shed light on an important yet elusive concept used in the blockchain literature as well as many other fields, and bring forward a data science pipeline which has been publicly shared and can be used to analyse the evolution of any concept in the scientific literature.

*A user to user network of interactions exists around Dark Web Marketplaces, which is resilient to their closure and is worth even more than the whole DWMs ecosystem.*

In chapter 7, we have studied the network of user to user interactions between users of Dark Web Marketplaces. While episodic proofs of its existence had been hinted at by small scale interview-based studies [15, 301], we reveal its existence and large size thanks to blockchain records of Bitcoin transactions. First, we characterize the network, showing how it generates a larger trade volume than the whole DWMs ecosystem itself. Then, we used a statistical temporal network based method [1, 28] to identify stable user to user pairs, forming the backbone of the trading network. We show how the users participating in these pairs trade more and generate more volume than other users. By analysing their temporal patterns, we show how DWMs have likely acted as a meeting point for a relevant fraction of them, and finally show their resilience to the closure of the DWM that spurred their formation. Our results reveal the existence of a relevant phenomenon so far ignored by the scientific literature and institutions alike, whose decentralized collective behaviour is of interest to complex system scientists and economists, as well as police authorities. Importantly, this phenomenon could be studied only because the blockchain provides large scale interactions data.

*Heterogeneous rarity patterns drive price and trade in the collectible NFT market.*

In chapter 8, we have investigated the dynamics of collectible NFT collections, the largest NFT market existing at the time of publication of this thesis. First, we measure the within-collection rarity of NFTs based on human readable attributes, showing how most collections are characterized by heterogeneous rarity patterns. Then, we study how rarity impacts the market performance, showing how, on average, rarer NFTs sell for higher prices, are traded less and lead to higher returns on investment with less associated risk. Our results provide some first hints at drivers behind the dynamics and success of this new and growing blockchain-based ecosystem.

These results shed light on multiple aspects of different blockchain-based ecosystems. Yet, they also open up new avenues for research and stimulate multiple follow-up questions. Here, we list some of the topics which we plan to investigate in the immediate future.

*Can we reliably identify buyers and sellers from DWM blockchain transaction data?*

The Bitcoin transactions dataset used in this thesis comes with a strong limitation: money is either going into or coming out of the market entity, hiding the buyer-seller link. Importantly, a buyer might receive money from the market as they might be taking some money out of their balance, similarly to a seller depositing some money on its market balance, making it hard to simply identify buyers and sellers, and limiting the kinds of analysis which can be reliably done. Research should work on possible heuristics or machine learning based methods to identify buyers and sellers. Their identification would improve existing analysis and would lead to more insights on their behaviour on DWMs and more in general in their Bitcoin transactions. For instance, our dataset contains their whole transaction history, allowing for a more general description of their behaviour.

*What are the main properties of NFT transactions networks?*

Our analysis has focused on a specific subset of NFT collections, and the role of rarity in driving their performance. However, future analysis could move the focus from NFTs to NFT traders, analysing their network of interactions. For example: do buyers specialize in specific kinds of collections? Are there more central or important actors in the network? Is it connected? Are there more important communities? Answering to these questions would provide important insights into the system dynamics, and provide important information to participating actors and regulators alike. Moreover, such data is usually not available on similar traditional markets, and researching these questions would therefore potentially provide indications for these systems too.

*Extending the NFT market analysis to emerging visual properties*

Our work on NFTs has relied on human-readable visual attributes which were embedded into the NFT metadata by the creators themselves. For instance, in the case of Cryptopunks, one has the visual attributes beard, hat or alien. However, these traits are defined by the creators within a single collection, and are limited to what was originally and explicitly planned by the creators for this specific kind of collectibles. Relying on state of the art computer vision methods, one could extract new visual attributes which characterize the NFTs but were not explicitly encoded by the creators. These new traits could theoretically encompass multiple collections, and would even facilitate a longitudinal analysis of the evolution of NFTs across time in terms of their visual aesthetics.

*Can we use DWM transaction data to fit and benchmark microscopic economics models?*

In chapter 5 we have used Bitcoin transaction data and regulated markets transactions to fit and benchmark a network formation model to explain the macroscopic properties of buyer-seller networks. However, economic models can potentially provide more insights

into the behaviour of users, having their foundation on game theory, incentives and other microeconomic mechanisms. A natural line of research is therefore to study such models, and combine their insights with the datasets employed in this study, which are unique and generally hard to find for academic studies.

# Chapter A

## Appendix to chapter 3

### A.1 Data pre-processing

In the following, we describe the DWMs dataset in more details, by focusing on how listings were stored and how we formed the COVID-19 categories in Table 3.1, that is, *PPE*, *medicines*, *guides on scamming*, *web domains*, *medical frauds*, *tests*, *fake medical records*, *ventilators*, and *COVID-19 mentions*.

Table A.1: **Selected attributes of the listings under consideration**, along with a brief explanation of their respective purposes.

Attribute of a listing	Explanation
“Listing body”	Description of the listing as it appears in the DWM
“Listing title”	Title of the listing as it appears in the DWM
“Marketplace name”	Name of the DWM
“Shipping information”	Where the listing is declared to ship from and to
“Time”	When the listing is observed
“Quantity”	Quantity of the listing sold
“Price”	Price of listing
“Vendor”	Unique identifier of the vendor

The listings appearing on the DWMs were crawled and stored according to selected attributes. While a brief explanation of these attributes is already presented in Table A.1, here we focus on those attributes which involved some pre-processing before the analysis, that is, “Shipping information,” “Quantity,” and “Price.” The “Shipping information” attribute was initially stored considering what the vendor declared. Then, it was standardised among vendors to correct any misspellings, using the standard python library *pycountry*. Vendors may declare a specific country, like United States, a continent, like Europe, or the entire world, which we standardise here as worldwide. The “Quantity” attribute was instead re-

trieved from the title of the listing using Facebook open-source library *Duckling* [350], then it was manually checked and corrected during an annotation process. The “Price” attribute on DWMs was displayed in the listings in various currencies, such as cryptocurrencies and fiat currencies. In order to standardise and properly compare listing prices, we converted prices to USD at the daily conversion rate. Rates were taken from Cryptocompare [351] for cryptocurrencies, and from the European Central Bank [352] for fiat currencies.

The attributes “Listing body” and “Listing title” in Table A.1, representing the title and description of the listings, were used to select the COVID-19 categories in Table 3.1. To this end, we prepared two sets of keywords as shown in Table A.2. Every selected COVID-19 listing contained either a word in the “Listing body” that matched one keyword in the first set or a word in the “Listing title” that matched one keyword in the second set. The rationale behind this choice was that the listing title was usually more precise on the product sold, whereas the body might contain promotions of other items the vendor was selling in other listings. At the same time, the vendor might mention COVID-19 in the body for various reasons, which we analysed in the main text. In order to classify listings in either COVID-19 specific listings (that is, *PPE, medicines, guides on scamming, web domains, medical frauds, tests, fake medical records, ventilators*) or COVID-19 mentions, we ran a regex query in google *bigquery*. We remark that the chosen method returned words containing a string equal to one of our keywords. For instance, with the keyword chloroquin, we detected also chloroquine and hydroxychloroquine. After this automatic filtering step, we manually checked the selected COVID-19 related listings to further improve the accuracy of our sample. In order to minimize human error, at least two authors from publication [I] checked each of these listings. A limitation of our approach was that keywords considered were in English. Therefore, even if drug names such as chloroquine were common to many languages and we detected some listings in a non-English language, our sample of COVID-19 related listings was biased toward the English language.

While each listing had an associated url to determine its uniqueness, which allowed us to track listing over time, vendors receiving bad reviews sometimes put identical copies of the same listing online. To overcome this issue and correctly count the number of listings, we determined a listing as unique if it had the combination of “Listing body,” “Listing title,” “Marketplace name,” and “Vendor” different than any other listings. For instance, if two listings had the same title and body but were sold in two different DWMs, we considered them as two different listings. Also, we considered at most one observation for each unique listing per day. The total number of unique listings and observations of these listings in each DWM is available in Table 3.2.

Table A.2: **Keywords used to sample COVID-19 specific listings from the DWMs in Table 3.2.**

First set of keywords checked against the words included in the attribute “Listing body” in Table A.1
corona virus, coronavirus, covid, covid-19, covid19
Second set of keywords checked against the words included in the attribute “Listing title” in Table A.1
anakinra, antidote, antiviral, azithromycin, baloxavir, baricitinib, bemcentinib, chloroquin, corona virus, coronavirus, covid, covid-19, covid19, darunavir, dexamethason, diagnosis, diagnostic, favipiravir, ganciclovir, glove, gown, lopinavir, marboxil, mask, n95, n99, oseltamivir, prevention, remdesivir, repurposed, ribavirin, ritonavir, sanitiser, sanitizer, sarilumab, siltuximab, surgical, thermo scanner, thermo-scanner, thermometr, thermoscanner, tocilizumab, umifenovir, vaccine, ventilator, ciprofloxacin, doxycyclin, metronidazol, amoxicillin

## A.2 Examples of listings related to COVID-19 in dark web marketplaces

Here, we present detailed examples of the selected listings. We consider both COVID-19 specific listings and COVID-19 *mentions*.

### A.2.1 COVID-19 specific listings

The most popular category of COVID-19 specific listings was *PPE*, which included mainly face masks. We detected listings selling small quantities of masks, like “KN95 Face Mask for Corona Virus box of 50” priced at 50 USD, while others proposed wholesale deals, as in “AFFORDABLE 20 BOXES OF SURGICAL FACE MASK (WHOLESALE PRICE)” in which 5000 masks were available at 2,000 USD.

The second most popular COVID-19 category was *guides on scamming* and includes listings explaining how to stole several kinds of COVID-19 related relief funds. Specifically, a subset of these listings were about the Small Business Administration loan in the USA. They provided step-by-step instructions, with constant updates to ensure the scam activities were effective. One listing in particular suggested: “I do not recommend taking more than 10,000 of the approved amount, because after that cashing out becomes a little harder.” The price of this listing was 113 USD.

MagBo was the only DWM selling listings in the *web domains* category. These listings may cause a potential threat to public health. They may be the actor of several phishing activities or sell scams. Examples of these web domains were “coronavintheworld.com,” “covid-conspiracy.net,” and “coronavirusmasks.in.” Prices of these domains were low and

less than 10 USD.

Listings on the *medicines* category were composed mostly by chloroquine, hydroxychloroquine, and azithromycin. We registered several wholesale deals, as in “9000 tabs hydroxychloroquine 200mg (USA AND CANADA ONLY)” where 9,000 tabs were sold for 1,194 USD. The smallest quantity we detected was 50 pills “chloroquine 50pills for 250\$,” sold at 250 USD. We also noticed that vendors often specified the size of the pill, being it 200mg, 250mg, or 500mg. The azithromycin was usually sold together with hydroxychloroquine as a prescription against COVID-19. One example of it was “hydroxychloroquine sulfate 200mg and azithromycin 250mg,” where an unknown quantity of these drugs was sold for 40 USD.

In the COVID-19 category of *medical frauds*, the most prominent listings were vaccines. Despite at the moment of writing of the first version of this manuscript (July 2020), vaccines are far from being actually developed, they were sold in DWMs since March. These listings included both low price vaccines like “complete order free shipment COVID19 VACCINE,” sold at just 200 USD, or high price one like “Covid-19 Vaccine. Lets keep it low key for now,” priced at 15,000 USD. In addition, among the listings in the *medical frauds* category, one could find potentially dangerous illicit drug mixes with claimed curative power against COVID-19, like “Protect yourself from the corona virus:” a marijuana based drug mix supposedly helpful in recovery from coronavirus infection. Other *medical frauds* included a 300 USD “CORONAVIRUS DETECTOR DEVICE, SAVE LIVES NOW” or a 1,000 USD “Buy CORONAVIRUS THERMO METER.”

*Tests* category of COVID-19 specific listings count a few different items. We detected listings in the *tests* category both at low quantities, such as, “25 pcs COVID-19 (coronavirus) quick test,” sold for 430 USD, or at very large one, like “Corona Virus Test / Covid-19 Test Kits ( 5000Pcs),” for a price of 7,500 USD.

The three listings in the *fake medical records* categories can be used to fake COVID-19 diagnosis. One of these listings said in its title: “Novelty/Fake Medical Records! Any diagnosis, custom made.” And in its body claimed “The right medical excuse can get you out of anything, and open many doors,” with a following disclaimer “IT IS UP TO YOU TO USE THESE ETHICALLY AND LEGALLY!” The price for this listing was 20 USD only, which could favour its wide adoption.

The two listings in the *ventilator* category were ICU ventilators. They were advertising fundamental hospital instrument, such as, “ICU Respiratory Ventilators , Emergency Room Vents” sold at 800 USD or “BiPAP oxygen concentrator ventilato Amid Covid-19” for 2,000 USD.

### A.2.2 COVID-19 mentions

We describe three examples of listings in the COVID-19 *mentions* category. The listing with title “Best Organic Virginia Bright Tobacco Premium quality 600g” refers to the lockdown in its body as “unfortunately we have to respect coronavirus lockdowns, in order to ensure as

much security as possible, we had to choose one type of shipping that is unfortunately much more expensive while lockdowns last.” Another listing with title “(Out of Stock! Lower Price for Pre-orders Only) Testosterone Enanthate 250mg/ml - 10ml - Buy 4 Get 1,” mentions in the body that they “are currently out of stock of this product due to our oil suppliers not being able to get their raw powders shipped to them because of the Coronavirus” and they “have lowered the price a little to help make up for this delay.” A third listing mention a sale directly in the title “COVID-19 SPECIAL OFFER 1GR CROWN BOLIVIAN COCAINE 90% £65,” and link the discount with the distress caused by the pandemic.

### A.3 Timeline of the COVID-19 pandemic

In this Section we aim at providing a summary of the main events related to the pandemic, focusing on the ones cited in the main text and listed in Table A.3. This is by no means a complete summary of the COVID-19 pandemic timeline.

The first event to gain international attention and make the public aware of the coronavirus was the decision from China to lockdown the city of Wuhan, first epicenter of the pandemic, on January 23, 2020 [125]. The virus then found its way to Europe, where the first country to be heavily hit by the pandemic was Italy. The Italian government decided to lockdown the entire country on March 9, 2020 [149]. The virus rapidly spread in Europe and internationally, with cases appearing more and more in the United States, leading USA’s President Donald Trump to first take a stance on the possibility of using chloroquine to cure individuals infected from COVID on the March 18, 2020 [154]. The epidemic started to heavily hit the United States and cases were surging almost everywhere in the world: 70 days after the lockdown of Wuhan, the worldwide count of infections had already surpassed 1 Million cases on April 3, 2020 [165]. On March 27, President Trump signed the Cares Act with the first economic aids to whose affected by COVID-19 [155]. After that, he explicitly promoted the use of hydroxychloroquine on April 5, 2020, before any official medical trial ended [162]. In April the situation started to become asymmetric. In Europe, thanks to the many policies in place, the COVID-19 became less threatening [161] and lockdowns started to be eased [151]. USA and other countries were instead seeing a rise in cases, and the USA Senate prolonged the small business rescue fund [156].

In May, President Trump declared he was now taking Hydroxychloroquine preventively against COVID-19 [163], while in July, he posted a video (label banned by Twitter) diffusing misinformation about the medicine [164]. The second wave of contagions hit Spain in September, and few weeks later the entire Europe [152], while The USA saw the failing on the negotiations around a second relief package [157]. Several new lockdown measures have took place in November in Europe [153] and, through that month, the number of COVID-19 related new infections has started to reduce. In the meantime, the USA were continuing to register a high number of new contagions.

Table A.3: **Significant COVID-19 events.** We defined an acronym for each event and reported it in the main text plots. Please note that this list does not intend to be exhaustive or to establish a ranking between events.

Date	Event	Acronym
2020-1-23	Wuhan Lockdown [125]	Wuhan
2020-3-9	Italy Lockdown [149]	Italy
2020-3-18	USA’s President Trump first refers to chloroquine [154]	Trump 1
2020-3-27	USA’s President Trump signs the CARES act [155]	Cares
2020-4-3	1M COVID-19 cases worldwide [165]	1M cases
2020-4-5	USA’s President Trump promotes the use of chloroquine and hydroxychloroquine against COVID-19 [162]	Trump 2
2020-4-24	COVID-19 cases in Europe are beginning to slow down [161]	Europe
2020-5-18	USA’s President Trump declares he is taking hydroxychloroquine preventively against COVID-19 [163]	Trump 3
2020-6-9	Governments start to lift lockdown measures around the world [151]	Openings
2020-6-30	USA senate agrees to extend small business rescue [156]	SBA
2020-7-28	Twitter limits Donald Trump Jr’s account for posting COVID-19 misinformation [164]	Trump 4
2020-9-7	Spain is the first country in Europe to record half a million COVID-19 cases [152]	Second wave
2020-9-10	Negotiations for the Heroes act keep failing [157]	Heroes
2020-10-31	PM announces four-week England lockdown [153]	New lockdowns

## A.4 Additional material

In this Section we provide additional material that support our main findings. In Table A.4 we provide more details on the 30 DWMs considered in our study. In particular we indicate the main specialization of the DWMs, i.e., the main category of products sold. If it is “Mixed”, it means that the DWM is not specialised in any particular category of goods. In the description we instead put information on the DWMs, with more details where available. All this information has been researched and compiled by the authors, with particular help given by Flashpoint Intelligence [30].

In Table A.5 we provide a Table reporting the different COVID-19 related medicines which were found in the listings. The medicines were selected as they have been found or claimed to be effective against COVID-19 [138]. The number of listings related to each medicine is also reported, noting that some listings sell more than one medicine (e.g. listings selling both hydroxychloroquine and azithromycin).

In Figure A.1 we plot the distribution of listings per vendor in log-log plot, showing a clear power-law shape with exponent -2.0. In the inset of Figure A.1, we show the histogram using linear spacing, through which we understand that most vendors sold very few COVID-

19 specific listings, while few vendors going as high as 91 different listings. We noted that 80% of the vendors had indeed less or equal than 5 listings.

In the main text, we performed a longitudinal analysis of the time evolution of all COVID-19 specific listings and all listings mentioning COVID-19, as well as the *PPE* and *medicines* categories, as shown in Figure 3.3. Now, we provide a similar analysis for COVID-19 *guides on scamming*, as illustrated in Figure A.2. We observe they first appeared in March, when the first lockdown measures were adopted. The number of listings then started increasing after the Cares act was introduced in USA [155]. Other peaks coincide to the extension of the SBA loan program in July [156] and to the failing of negotiations on the Heroes act [157], after which the number of listings decreased up to April levels. Listings in the *guides on scamming* category teach people how to take advantage of several kinds of COVID-19 relief funds, or other pandemic related scam opportunities. In many western countries, new relief funds were signed on a monthly basis constant updates made on the relative listings on DWMs.

In order to complement Figure 3.4(a) and (b) in the main text and properly show the peaks of Wikipedia page visits and tweets, we create Figure A.3. The new representation of Figure 3.4 does not modify the claims made in the main text and how major event related with COVID-19 impacted public attention.




Figure A.1: **Probability distribution function (Pdf) for the number of listings per vendor.** The power law fit results in an exponent of -2.3. In inset, the histogram of the number of listings per vendor, with a vertical line showing the 80<sup>th</sup> percentile.




Figure A.2: **Time evolution of the active COVID-19 specific listings in the *guides on scamming* category.** (a) Seven-days rolling average of these observed listings at a given time. Black dashed vertical lines corresponded to significant COVID-19 world events, see Appendix A.3. (b) Seven-days median price with 95% confidence interval for these observed listings.




Figure A.3: **Wikipedia page visits for pages relative to (a) *PPE*, (b) hydroxychloroquine, chloroquine and azithromycin.** Number of tweets mentioning (c) *PPE*, (d) hydroxychloroquine, chloroquine and azithromycin. Panels (a) and (b) corresponds to Figure 3.4(a) in the main text, while panels (c) and (d) to Figure 3.4(b). The main difference between these panels and Figure 3.4(a) and (b) is the linear scale on y axis.

Table A.4: **List of all DWMs, together with their specialization and a brief description.**

DWM	Specialization	Description
Atshop	Digital Goods	Atshop e-commerce marketplace platform
Black Market Guns	Weapons	Weapons Marketplace, now exit scammed according to onion.live [73]
CanadaHQ	Mixed	Multivendor cryptocurrency marketplace
Cannabay	Drugs	Russian language drug marketplace focusing on cannabis
Cannazon	Drugs (Cannabis)	Drug marketplace for cannabis products only
Connect	Mixed	Social network hosting a marketplace selling illicit goods
Cypher	Mixed	Multivendor market selling drugs and digital goods.
DarkBay/DBay	Mixed	Multivendor cryptocurrency DWM selling digital goods, drugs, and services
Dark Market	Mixed	Multivendor cryptocurrency DWM selling digital goods, drugs, and services
Darkseid	Weapons	Weapons DWM
ElHerbolario	Drugs	Single-vendor shop, selling just 3 products, primarily leaning towards Cannabis
Empire	Mixed	Alphabay-style DWM with BTC, LTC, XMR, MultiSig, and PGP 2FA
Exchange	Mixed	Chinese language marketplace
Genesis	Digital goods	Marketplace selling digital identities for account takeover activities
Hydra	Drugs	Russian language drug DWM
MagBO	Digital Goods	Shell, account and card shop
MEGA Darknet	Mixed	Russian language DWM
Monopoly	Drugs	Multivendor market that is primarily focused on drugs
Mouse In Box	Digital Goods	Marketplace that sells packages of login and session information acquired from web browsers with a stealer malware.
Plati.Market	Digital goods	digital goods DWM
Rocketr	Digital goods	Marketplace for the sale of illicit digital goods
Selly	Digital goods	Marketplace for the sale of illicit digital goods
Shopyy.gg	Digital goods	Marketplace for the sale of illicit digital goods
Skimmer Device	Skimmer devices	Marketplace selling skimmer devices
Tor Market	Drugs	Drug DWM focused on supplying the drug marketplace in New Zealand
Torrez	Mixed	Multivendor market using wallet-less payments.
Venus Anonymous	Mixed	Multivendor DWM selling digital goods and drugs
White House	Mixed	Multivendor cryptocurrency DWM
Wilhaben	Mixed	German language DWM for the selling illicit goods
Yellow Brick	Mixed	Multivendor cryptocurrency DWM

Table A.5: **COVID-19 related medicines appearing in the listings**, together with a brief description and the number of listings related to that drug.

Medicine	Description	Listings
Hydroxychloroquine	Malaria medication	65
Azithromycin	Antibiotic often paired with hydroxychloroquine	51
Amoxicillin	Antibiotic medication	45
Chloroquine	Malaria medication	38
Ciprofloxacin	Antibiotic medication	6
Favipiravir	Antiviral medication used to treat influenza	5
Doxycycline	Antibiotic medication	4
Metronidazole	Antibiotic medication	4
Remdesivir	Antiviral medication	2
Lopinavir	Antiviral medication used to treat HIV	1

Table A.6: **COVID-19 related *guides on scamming* appearing in the listings**, together with a brief description and the number of listings related to that sub-category.

Topic	Description	Listings
SBA loan	how to illicitly get money from the USA Small Business Loan program [156]	19
Bank account	how to exploit pandemic related security to open bank accounts	16
Fraud Pack	pack containing multiple generic covid related frauds	7
Covid-19	Generic guides explaining how to exploit the pandemic in many different ways	7
Amazon	Amazon related fraud guides	6
GoFundMe	GoFundMe related fraud guides	4
Apple	Apple related fraud guides	3
Unemployment fund	How to illicitly get money from government unemployment funds	3
Other	Other COVID-19 related fraud guides	34

# Chapter B

## Appendix to chapter 4

### B.1 Examples of detected listings




Figure B.1: **Pfizer/BioNTech vaccine offered on Invictus.** Screenshots of a listing in the *approved vaccines* category offering the Pfizer/BioNTech vaccine at \$500 on the Invictus marketplace. We removed the contact information of the vendor, who invites the potential customer to have a direct contact. The screenshot was taken on February 6, 2021.

Table B.1: **Generic vaccine offered on a DWM.** Example of a listing in the *unspecified vaccines* category offering a generic vaccine, which does not specify the producer. Personal information of the vendor are hidden with # symbols.

Title	COVID-19 antidote from china. offering at 15k USD
Body	the covid-19 current massacre is supposed to have ended by now. while the who is trying to be selfish with human life, we are trying to save the lives. the real virus is the leaders. this vaccine should be used just once on one person and basically the giveaway price i put here is nothing compare your life. get your vaccine now in time. you can buy from me and resell at your price. contact me for more details. email: ##### wickr: ##### telegram: ##### kik: #####
Price	15,000 USD
Shipping from/to	N.A.
Vendor	#####
DWM	DarkBay




Figure B.2: **Proof of vaccination offered on Hydra.** Proof of vaccination offered at 4,000 rubles (55 USD) detected on the Hydra marketplace. The original language of this listing is Russian and its translation in the English language is given in Table B.2. The screenshots were taken on February 6, 2021.

Table B.2: **Translation of the proof of vaccination offered on a Hydra.** English translation of the listing in Fig B.2(b). We use Google Translator to translate the text from Russian to English.

Title	Coronavirus / Covid / Corona vaccination certificates
Body	Hello Friend! The coronavirus has made adjustments to everyone's life. And more than once many of us will face its consequences. One of them is the provision of a COVID-19 vaccination certificate, which can be presented at work, at the place of study and in any other place where it is required. To order, you need to indicate your full name, date of birth, date of the issued document. If you need a non-Moscow institution, you will need to pay extra for the production of the necessary seals. Production time is 2-7 days. Delivery when sent by courier will have to be paid for upon receipt. It is possible to send by registered or regular mail, then we will take on this heavy burden. An ordinary letter has no track, therefore, until we receive the letter, we will remain in the dark about its fate.
Price	55 USD
Shipping from/to	Russia and neighbouring Eastern countries/Russia and neighbouring Eastern countries
Vendor	#####
DWM	Hydra

Table B.3: Example of drug listing mentioning COVID-19, and problems with the United States Postal Services (USPS).

Title	(10) 30mg adderall pressed pills: us - us
Body	<p>Adderall is used in the treatment of attention deficit hyperactivity disorder (ADHD) and narcolepsy. It is also used as an athletic performance enhancer, cognitive enhancer, appetite suppressant, and recreationally as an aphrodisiac and euphoriant. It is a central nervous system (CNS) stimulant of the Phenethylamine class.</p> <p>You MUST Be A Minimum Of 18 Years of Age</p> <p>For Research Purposes Only Not For Human Consumption Refund Policy ALL SALES ARE FINAL!</p> <p>REFUNDS WILL NO LONGER BE ISSUED DUE TO SCAMMING! RESHIPS ARE ALWAYS AVAILABLE ON A CASE BY CASE BASIS, AND USUALLY ONLY WHEN A TRACKING NUMBER NEVER ORIGINALLY SCANS, OR GETS STUCK FOR 15+ DAYS.</p> <p>NO reships will be sent in the event of a tracking status of RETURN TO SENDER or UNDELIVERABLE AS ADDRESSED.</p> <p>Reships DO NOT qualify if a package status is marked as Delivered or indicates the package is In Transit to its destination. If the package is in the system, please wait it out for the package to arrive. THE USPS IS UNDERFUNDED AND MAY BECOME UNRELIABLE COMPARED TO THE PAST! (ESPECIALLY DURING COVID-19 AND HOLIDAYS!)</p> <p>Use a real name and address for your package. If a package is stuck IN TRANSIT for a few days, and a tracking number is given to you, please call USPS to locate it. THOUGH I MAY CARRY NON SCHEDULED RESEARCH CHEMICALS, DO NOT CLAIM TO KNOW THE CONTENTS. NO PACKAGES WILL EVER REQUIRE A SIGNATURE!</p>
Price	59.13 USD
Shipping from/to	USA/USA
Vendor	#####
DWM	Dark0de Reborn

Table B.4: **Example of drug listing mentioning COVID-19, and ensuring safety measures are taken.**

Title	black diamond og sfv og shake popcorn the
Body	<p>thank you so much for shopping with us we are confident you'll love your order while your here take a moment to browse through our vendor page to see the many great strains bulk orders and emeraldgallipot promotional offers we have to offer what we offer fast communication all msg are answered within hrs fast delivery product will be shipped on the next business day after order confirmation via usps priority mail stealth packaging vacuum sealed odorless sterile packaging package tracking available upon request three days after order confirmation full refund replacement if tracking confirms package seized lost</p> <p>what we ask please provide your full address immediately in pgp format in buyer's note use your full name and double check your address deliveries that tracking confirms lost because of errors in provided information are not available for refund or replacement</p> <p>all shipping addresses must be in the following format name john doe address nameless ln city state zip city xx</p> <p>please finalize asap upon receiving package please leave a positive rating if you are unhappy with your order please tell us we are happy to work with you to satisfy your needs</p> <p>strain highlights black diamond og indica dominant hybrid backberry kush diamond og the</p> <p>flavor aroma a cross between blackberry and diamond og its flowers have a glittery trichome covering and purple coloring that make it a beautiful gem to look at the strains aroma is musky and earthy almost like a deep red wine</p> <p>euphoric effects black diamond is known to cause fits of giggles and is a great strain for hanging out with friends and creative expression</p> <p>medical benefits ideal for patients who need strong medication but still want to be active and sociable this strain tends to make consumers extremely hungry making it a good choice for those looking to increase their appetite just make sure you have some snacks on hand</p> <p>san fernando valley og sativa og kush direct the</p> <p>flavor aroma sfv og by cali connection is a sativa dominant hybrid that is as the name indicates this og kush relative originates from californias san fernando valley although their names are barely distinguishable sfv og kush is actually the afghani crossed child to sfv og leading with aromatic notes of earthy pine and lemon</p> <p>euphoric effects creates a long lasting head haze and full body effect that leaves you feeling happy and relaxed without damping your energy</p> <p>medical benefits great for patients who need strong pain relief but dont want to be stuck on the couch</p> <p>note we here at the emeraldgallipot take our customers safety as our highest priority and to help protect you against the spread of the coronavirus all packages we send are being thoroughly sterilized with a mild disinfectant and bleach solution prior to shipping for your protection stay safe out there</p>
Price	50 USD
Shipping from/to	USA/USA
Vendor	#####
DWM	Torrez

Table B.5: Example of drug listing mentioning COVID-19, and “stocks are almost exhausted by Corona”.

Title	grams speed paste normal quality
Body	<p>When you place an order you agree with our conditions!  Offer: 5 Grams Speed Paste Normal Quality  This product is made from high grade washed A-Oil  Purity: 45% up to 55%  Approximately 20% of the weight is lost during the drying process  For any questions feel free to contact us, we are happy to help you!</p> <hr/> <p>Welcome to ##### The best speed (amphetamine) products on the market! We sell from the normal quality till the highest quality you can get! Our sending fits every mailbox! We ship from Monday till Friday!  We ship from Germany and we know how to ship! It is important for us that all orders arrive in all safety!  SHIPPING TIME Europe: 2 to 7 Business days Worldwide: 4 to 20 Business days  REFUND and RESHIP If orders not arrive please send us a message and we find a solution. In case of non-arrival, we will reship 50% or a 50% refund. Mistakes made in the address-format we will never reship or refund. New buyers without any order history we never refund or reship.  Please give us some great feedback if you are happy with us!  AmphetamineCowboys</p> <hr/> <p style="text-align: right;">UPDATE 13-02-2021</p> <p>Dear customers, From today 13-02-2021 we will go into vacation mode for 10 days until 23-02-2021. We do this because we have a lot of money in escrow and our stocks are almost exhausted by Corona Covid 19. New stocks are on the way but unfortunately it is slowing down due to Covid bullshit. We do not want to disappoint. We will receive new stocks next week so that we can continue on 23-02-2021. Of course all accepted orders have been shipped including today! We are online every day for all your questions about the shipped orders or for any other questions. Hoping for some understanding from you, we will be back soon on 23-02-2021. All be safe and hope to see you soon!</p> <p>Sehr geehrte Kunden, Ab heute 13.02.2021 werden wir fr 10 Tage bis zum 23.02.2021 in den Urlaubsmodus wechseln. Wir tun dies, weil wir viel Geld im Treuhandkonto haben und unsere Aktien von Corona Covid 19 fast erschpft sind. Neue Aktien sind auf dem Weg, aber leider verlangsamte sie sich aufgrund von Covid-Bullshit. Wir wollen nicht enttuschen. Wir werden nchste Woche neue Aktien erhalten, damit wir am 23.02.2021 weitermachen knnen. Natrlich wurden alle angenommenen Bestellungen auch heute noch versendet! Wir sind jeden Tag online fr alle Ihre Fragen zu den versendeten Bestellungen oder fr andere Fragen. In der Hoffnung auf ein Verstndnis von Ihnen werden wir bald am 23.02.2021 zurck sein. Alle sind in Sicherheit und hoffen, Sie bald zu sehen!</p>
Price	17 USD
Shipping from/to	Germany/Worldwide
Vendor	#####
DWM	White House

## B.2 DWMs offering COVID-19 vaccines

Table B.6: List of DWMs analysed.

Type of products	DWM
COVID-19 vaccines	Agartha, Asap, Babylon, Bigblue, Cypher, Dark fox, Hydra, Invictus, Kilos, Liberty, Mgm grand, Recon, Royal, Televend, The Canadian Headquarters, Torrez, World market, Yakuza, Yukon
COVID-19 related products	0day.today, Agartha, Asap, Bigblue, Corona, Cypher, Dark fox, Dark0de reborn, Darkmarket, Incognito, Kilos, Liberty, Magbo, Recon, Televend, The canadian headquarters, Torrez, Versus, White house, Yakuza
COVID-19 mentions	0day.today, 24HoursPPC, ASAP, Agartha, Amigos, Apollon Marketplace, Asean, Atshop, Auction DB, Aurora, Babylon, Big Brother House, BigBlue, Blackhole, CannaHome, Cannabay, Cannazon, Cartel, Cindicator, Connect, Corona, Cypher, Dark Fox, Dark Leak Market, Dark0de Reborn, DarkBay/DBay, DarkMarket, Database, Deep Sea, Deepsy, DutchDrugz, Empire Market, Exchange, FSpros, Faceless, Flugsvamp 3.0, Fullzbuy, Genesis marketplace, HeinekenExpress, Hexablaze, Hookshop, Hydra, Incognito, Invictus, Kilos, Liberty, MEGA Darknet Market, MGM Grand, MagBO, Market Deepmix, Metropolis, Monopoly, Mouse In Box, Namaste LSD, Olux, Opiate Connect, Pentagon, Plati.market, Psylab Seeds, RNJLogs, Recon Search Engine, Royal, Russian Market, SEOclerks, Scans24, Sellix, Shoppy.gg, Silk Road 3.1, Silk Road 4, Tea Horse Road, Televend, The Canadian HeadQuarters, Tor Market, Torrez, UAS, Versus, WTN Market, White House, Willhaben, World Market, Xleet, Yakuza, Yellow Brick marketplace, Yukon

Table B.7: **Vaccine listings detected on DWMs.** Some vendors and DWMs offer vaccines that belong in more than one category.

Category	Unique listings	Vendors	DWMs
<i>Unspecified vaccines</i>	94	61	13
<i>Approved vaccines</i>	74	44	7
<i>Proofs of vaccination</i>	80	42	10
Total	248	134	19

### B.3 Additional material




Figure B.3: **Summary of key events related to DWMs and COVID-19.** Availability of listings offering vaccines on dark web marketplaces (top), together with main COVID-19 related events of the vaccination campaign (bottom).




Figure B.4: **Lifetime of listings on a DWM.** Number of days during which listings were active on a DWM.




Figure B.5: **Categories of vaccines offered on DWMs.** Number of DWMs offering a vaccine in a given category. Only the lower triangle of the matrix is shown because it is symmetric, where its diagonal represents vendors offering only listings in that category.




Figure B.6: **Vendor statistics.** Histograms representing the number of vendors offering a certain amount of vaccines listings, in panel (a), and the number of vendors trading in a given amount of DWMs, in panel (b).




Figure B.7: **Price of COVID-19 approved vaccines.** Boxplots of the prices in USD at which vaccines were offered. (a) Price of listings in the three categories considered. (b) Focus on the listings offering approved vaccines. “J&J” stands for Johnson&Johnson. Horizontal lines represent the median value, box ends the first and third quartiles, and whiskers minimum and maximum values, respectively.




Figure B.8: **Temporal evolution COVID-19 approved vaccines.** (a) Cumulative number of listings over time. (b) Average price over time, computed with a 90-days moving window. “J&J” stands for Johnson&Johnson.

Table B.8: **COVID-19 related products offered on DWMs.** Availability of COVID-19 related products since November 2020.

Category	Unique listings	Observations	Median price [USD]	Vendors	DWMs
Guides on scamming	50	885	50	36	15
Malware	4	19	NaN	3	1
Medicines	40	367	38.00	27	13
PPE	6	36	15.00	3	3
Test	17	85	211.12	11	8
Web domain	38	184	4.00	13	1




Figure B.9: **Time evolution of products mentioning COVID-19.** Number of active listings in time in each product category, according to the clustering described in the main text.




Figure B.10: Time evolution of fraction of all listings mentioning COVID-19 related themes.

# Chapter C

## Appendix to chapter 5

### C.1 Computing the Memory Kernel

In this section we describe the computations made in order to estimate the buyer memory parameters  $c, \beta$ . As detailed in the main text, buyers are grouped in different classes according to their final degree (number of different sellers they purchased from) at the end of the considered periods. Classes are divided in powers of 2, for example the first class includes buyers with degree 1, the second with degree between 2 and 3, the third between 4 and 7 and so on. In our dataset, a buyer has a unique identifier in each market product, not allowing to follow their behaviour across different markets. In order to reduce the noise in the data, all markets are aggregated together for the following computations.

In order to estimate the memory parameters from eq.2, defined in the main text, the first step is to estimate the conditional probabilities  $P_k(n+1|n) = P_k(n)$  of buying from a new  $n+1^{th}$  seller when you already bought from  $n$  different ones. To do so, we count the number of buyers  $b_k(n)$  in class  $k$  who go from degree  $n$  to  $n+1$ , and we count the total number of purchases  $p_k(n)$  they made when they had degree  $n$ .

$$P_k(n) = b_k(n)/p_k(n) \tag{C.1}$$

In order to reduce the noise on the computation of  $P_k(n)$ , we limit the computation to  $n \leq k$ . This way, all buyers in degree class  $k$  go from degree  $n$  to  $n+1$ , as their final degree is at least equal to  $k$ . The numerator is therefore constant, and equal to  $N_k$ , the number of buyers in degree class  $k$ . Equation  $P(n)$  then reads:

$$P_k(n) = N_k/p_k(n) \tag{C.2}$$

Assuming that for a given degree  $n$  events are independent, or in other words that users

behave independently of each other, and checking that  $1 \ll N_k \ll e_k(n)$ , we can estimate the uncertainty of  $P_k(n)$  as follows:

$$\sigma(P_k(n)) = \sigma_k(n) = \sqrt{P_k(n)(1 - P_k(n))/e_k(n)} \quad (\text{C.3})$$

Having estimated the curve  $P_k(n)$  for each degree class  $k$ , we can fit eq.2 to each curve separately. To do so, we do a numerical least square optimization, estimating the values of  $\beta$  and  $c$  for each degree class. Results are shown in Tab. C.1.

## C.2 Sampling of Product Markets

The e-commerce platform contains data on 144 product markets. We sample 28 DWMs to fit with our model. The 28 product markets are sampled to ensure all products are represented. In particular, products can be grouped together in higher-level markets, from which we sample on product each. To make an example: our dataset may contain two product markets in the fruit group, namely apples and pears. In the sample for the model simulation we choose only one of the two, taking care that the 28 sampled product markets are representative of the heterogeneous market size of our dataset.

## C.3 Model parameters estimation

As detailed in the main text, we employ a data-driven approach to estimate the model parameters for each product market. The model parameters, described in the main text, are chosen in the following way. The number of time steps of the simulation  $T$  is chosen such that the average number of transactions in the simulation is equal to that of the data. Similarly, the number of buyers and the number of sellers are the same as in the market we are simulating. The users' activity is instead drawn from the empirical activity distribution. The memory parameters, namely  $c$  and  $\beta$ , are fixed to  $10^{-3}$  and  $10^{-1}$  respectively, values that well represent the different values fitted from the empirical dataset as shown in Table C.1. Finally, the preferential attachment parameter  $\mu$ , describing the increment of a seller's attractiveness after a sale, is estimated by Maximum Likelihood Estimation (MLE). To do so, we simulate the model for each value of  $\mu$  on a grid, ranging from 1 to 500, and then compute the associate negative log-likelihood computed comparing the data to the simulated attractiveness distribution, and choose the value minimizing the quantity. For instance, the likelihood of our empirical data is computed as the product of the probabilities that each data point had according to the attractiveness distribution built with the simulated data. We employ this simple approach to estimate  $\mu$ , by only analysing the attractiveness distribution, as the scope of this work is to study and reproduce stylized facts, and not to propose a detailed model precisely reproducing all details of a given product market. For this reason, even the value of  $\mu$  itself assumes relative importance, as its order of magnitude determines the agreement with the data, but small variations in the precise value are meaningless in the

context of this study. For completeness, in Tab. C.3 we show the fitted value of  $\mu$  for each product market.

## C.4 Model simulation

Each simulation is tuned to simulate one specific product market. We fix the agents population according to the data: number of sellers  $N$ , number of buyers  $M$ , and simulation total number of time steps  $T$ , to fix the average total number of transactions in our simulations as in the data:  $\langle a_i \cdot \Delta t \cdot T \cdot N \rangle = t$ , where  $a_i$  is the buyer activity as defined in the main text,  $\Delta t$  is the simulation time step (fixed to 1) and  $t$  is the total number of transactions present in the data. We realise 30 different realizations for each parameter set, and aggregate the final results.

## C.5 Additional Figures

In Fig. C.1, we show the histogram of the percentage of users with entropy zero doing just one transaction, in each product market. This percentage is always greater than 75%, but actually over 90% in most cases, showing how buyers with entropy zero can effectively be neglected when showing the buyer entropy distribution.

In Fig. C.2, we test the entropy distribution from Fig. 2a of the main article against a null model. In Fig. C.2a, we reproduce the buyer entropy distribution for each market (in red) and all markets (in yellow) for the e-commerce dataset. In Fig. C.2b, we show the same distributions for the same dataset reshuffling the transactions link, such that buyers maintain the same number of transactions but with randomly chosen sellers. The latter distributions show a narrower support with high values of entropy, showing that memory effects disappear in a null model preserving just the activity of buyers and sellers.

In Fig. C.3 to C.6 we show results of model simulation for 26 other product markets. The results show how the model is able to capture the main stylised facts of the buyer-seller network structure, with memory and preferential attachment both necessary to capture different aspects of the structure.

In Fig. C.7 to C.10 we show the temporal evolution of the top 50, 100 and 200 sellers degree distribution for 26 other product markets, represented as boxplot for 9 equally spaced time steps. The model is consistently able to reproduce the temporal evolution of the degree distribution, as shown by the cores of the boxplots (interquartile range) overlapping.

In Fig. C.11 we show the duration of each DWM in our dataset, color-coding by the average daily volume of transactions in USD. Our dataset covers all major DWMs from their onstart in 2011 with Silk Road Marketplace. DWMs are heterogeneous in daily volume, with some being just over our threshold of 20,000 USD, and Hydra Marketplace or AlphaBay Market close to 1M daily USD.

In Fig. C.12 we show the size of the 144 product markets in the regulated e-commerce platform dataset. In Fig. C.12(a), we show the total number of transactions in each product market, whereas in Fig. C.12(b) we show the total number of users (buyers and sellers). Product markets are heterogeneous in size, both w.r.t. number of transactions and number of users.




Figure C.1: **Most buyers with zero entropy only have done one transaction.** Histogram of percentage of buyers doing just one transaction for each e-commerce product market, among those with zero entropy. In most markets the percentage is well above 90%.




Figure C.2: **The entropy distribution is significant against a null model.** (a) Buyer entropy distribution for each market (red) and all markets (yellow) in the e-commerce datasets. (b) Same distributions build reshuffling the transaction links, such that buyers have the same number of transactions, but with random sellers. The distributions with the reshuffled data show considerably higher entropy, meaning that the buyers in the dataset show non-random effects of memory.




Figure C.3: **Model simulations for different markets - final distributions - Markets 3 to 9.** Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships.




Figure C.4: **Model simulations for different markets - final distributions - Markets 10 to 16.** Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships.




Figure C.5: **Model simulations for different markets - final distributions - Markets 17 to 23** Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships.




Figure C.6: **Model simulations for different markets - final distributions - Markets 24 to 28** Each row corresponds to a different market, whose simulations parameters are individually calibrated as detailed in the main text. From left to right, we show distributions for different quantities: attractiveness, seller degree, buyer degree, link weight and seller entropy. The comparison with the two model variations, without preferential attachment or without memory, shows the key role of both parameters in shaping the network: preferential attachment is crucial in reproducing highly active sellers, whereas buyer memory is fundamental to capture the heterogeneity of buyer-seller relationships.




Figure C.7: **Model simulations for different markets - temporal evolution - Markets 2 to 8** Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5<sup>th</sup> to 97.5<sup>th</sup> percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism.




Figure C.8: **Model simulations for different markets - temporal evolution - Markets 9 to 15** Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from  $2.5^{th}$  to  $97.5^{th}$  percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism.




Figure C.9: **Model simulations for different markets - temporal evolution - Markets 16 to 22** Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from  $2.5^{th}$  to  $97.5^{th}$  percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism.




Figure C.10: **Model simulations for different markets - temporal evolution - Markets 23 to 28** Each row represents one market. From left to right: temporal evolution of the degree distribution of the top 50 (left), 100 (center) and 200(right) sellers, representing the distribution at 9 equally spaced time steps with boxplots ranging from the first to the third quartiles, whiskers extending from 2.5<sup>th</sup> to 97.5<sup>th</sup> percentiles. The model better captures the temporal evolution of the top sellers degree for all product markets than the alternatives neglecting either the preferential attachment or the memory mechanism.




Figure C.11: **Dark Web Marketplaces** Duration of each market in our dataset, color coded by the number of transactions involving each marketplace. Each market is live at least for 180 days and averages at least 20'000 transactions per day.




Figure C.12: **E-Commerce Platform Markets** a): Plot of the number of transaction per each market of the e-commerce platform data. b): number of buyers and sellers in each market of the e-commerce platform data.

Table C.1: Memory kernel fitted coefficients  $\beta$  and  $c$  for every degree class  $k_{min} < \text{degree} < k_{min} * 2$

$k_{min}$	$\beta$	$c$
2	$0.0663 \pm 0.0003$	$0.00100 \pm 0.00004$
4	$0.0561 \pm 0.0003$	$0.00100 \pm 0.00004$
8	$0.0717 \pm 0.0003$	$0.0051 \pm 0.0001$
16	$0.0914 \pm 0.0004$	$0.0124 \pm 0.0003$
32	$0.1026 \pm 0.0004$	$0.0166 \pm 0.0005$
64	$0.1010 \pm 0.0006$	$0.0115 \pm 0.0005$
128	$0.0927 \pm 0.0008$	$0.0058 \pm 0.0005$

## C.6 Additional Tables

In Tab. C.1 we show the fitted coefficients of  $\beta$  and  $c$  for every buyer degree class. The coefficients are slightly different among the different degree classes, but of the same order of magnitude, allowing us to fix the value of  $\beta$  and  $c$  in the model simulation. This is only an approximation in the context of a model whose goal is to see the role of different mechanisms in determining the structure and evolution of the buyer-seller network. If the goal was to reproduce the finest detail of the network, or to make predictions on its evolution, we'd assign values of  $\beta$  and  $c$  according to the buyer degree class, where the degree is sampled from the data degree distribution.

In Tab. C.2 we show details on each DWM in the dataset. Details include the name, start and end data, reason of closure, type of goods traded, total number of transactions and total volume of transactions in USD.

In Tab. C.3 we show the value of the preferential attachment parameter  $\mu$  fitted for each product market. While values are heterogeneous, showing the different role of preferential attachment in each market, the precise value is not important in our study. Indeed, our only goal is to reproduce the main stylized facts of the data, not to reproduce the finest details of the network, and therefore changing the value of  $\mu$  around the fitted value would not change our conclusions.

Table C.2: **Details of the DWMs under study:** start, end, reason of closure, type, total number of transactions and total volume

Name	Start	End	Closure	Type	#Trx	Volume [USD]
Silk Road Marketplace	2011-01-31	2013-10-02	raided	mixed	840,987	131,604,274
Sheep Marketplace	2013-02-28	2013-11-29	scam	drugs	65,904	10,923,327
Black Bank Market	2013-10-18	2015-05-18	scam	mixed	89,444	13,152,830
Pandora OpenMarket	2013-10-20	2014-11-05	raided	drugs	73,127	8,401,191
Dream Market	2013-11-01	2019-04-30	voluntary	mixed	570,734	57,637,323
Silk Road 2 Market	2013-11-06	2014-11-05	raided	mixed	426,277	66,825,593
Agora Market	2013-12-03	2015-09-01	voluntary	mixed	911,094	141,473,388
Evolution Market	2014-01-14	2015-03-19	scam	drugs	372,822	47,578,872
Middle Earth Marketplace	2014-06-22	2015-11-04	scam	mixed	67,630	8,361,143
Russian Anonymous Marketplace	2014-08-29	2017-07-15	raided	mixed	1,109,126	80,478,841
AlphaBay Market	2014-09-01	2017-07-05	raided	mixed	4,263,740	546,010,808
Joker's Stash Market	2014-10-07	2021-02-03	closed	credits	998,687	153,138,403
Nucleus Market	2014-10-24	2016-04-13	scam	mixed	391,394	56,594,214
Abraxas Market	2014-12-24	2015-11-05	scam	drugs	168,642	21,854,042
Flugsvamp Market 2.0	2015-04-20	2018-10-02	closed	drugs	254,972	23,013,741
Unicc	2015-06-13	2021-02-10	active	credits	2,930,842	147,814,198
Hansa Market	2015-07-01	2017-07-20	raided	drugs	617,414	60,644,436
Hydra Marketplace	2015-11-25	2021-02-10	active	mixed	6,005,608	2,175,558,739
Wall Street Market	2016-09-09	2019-05-03	raided	mixed	681,825	48,153,667
TradeRoute Market	2016-11-22	2017-10-12	scam	mixed	137,722	16,969,504
Bypass Shop	2017-03-10	2020-12-27	closed	unknown	1,041,438	65,663,561
Buybest - GoldenShop	2017-11-13	2020-03-19	closed	unknown	386,046	24,449,110
Silk Road 3.1	2018-02-10	2019-08-27	scam	drugs	93,426	9,053,684
Empire Market	2018-04-01	2020-08-30	scam	mixed	454,473	154,457,692
Apollon Market	2018-05-03	2020-01-27	scam	drugs	106,395	12,902,953
Flugsvamp Market 3.0	2018-12-17	2021-02-10	active	unknown	291,018	39,344,294
FEShop	2019-08-14	2021-02-10	active	unknown	1,342,574	64,666,841
DarkMarket	2020-02-04	2021-01-12	raided	unknown	363,825	27,246,084

Table C.3: **Preferential attachment parameter  $\Delta$ .** Values of the preferential attachment parameter  $\Delta$  for each product market, fitted with maximum likelihood estimation on the attractiveness distribution.

Market	$\Delta$
1	85
2	21
3	155
4	225
5	400
6	220
7	90
8	14
9	50
10	27
11	290
12	15
13	13
14	70
15	33
16	14
17	175
18	140
19	240
20	210
21	180
22	200
23	65
24	35
25	190
26	410
27	220
28	230

# Chapter D

## Appendix to chapter 6

### D.1 Dataset

In this section we report additional information on the S2AG corpus and the *(de)centralization* dataset to complement the description in the article. In Table D.1 we report the number of publications in our dataset containing information about each attribute. In Table D.2 we report information on the publications' fields of study, as originally reported in the S2AG metadata. Each item can be classified in one or more fields. In the left table, we report all single fields or pairs that have more than 1000 documents in our dataset (papers can be classified in more than two fields, but the fields after the second are here discarded). In the right table we instead see how many papers have each of the 19 fields as first field in their classification, showing the predominance of Computer Science, and STEM subjects in general, in our dataset. Finally, in Fig. D.1 we show the raw number of papers and authors each year for the full S2AG corpus, compared with the *(de)centralization* dataset, showing how *(de)centralization* has a faster exponential growth than the general academic literature as indexed by semantic scholar, as the different exponent of the fitted exponential functions clearly show.

Attribute	#Publications
<i>title</i>	425,144
<i>paperAbstract</i>	396,201
<i>authors</i>	421,611
<i>year</i>	423,431
<i>inCitations</i> or <i>outCitations</i>	305,639
<i>fieldsOfStudy</i>	377,720
<i>doi</i>	253,464
<i>venue</i>	143,802
<i>journalName</i>	234,888

Table D.1: **Details of each attribute present in the S2AG dataset.** Number of papers (right column) in the *(de)centralization* dataset that contain information on the listed attributes (left column). The attribute "*inCitations*" and "*outCitations*" correspond respectively to received citations and given references.




Figure D.1: Comparison of the increase in number of papers **(a)** and authors **(b)** between the whole S2AG corpus and the *(de)centralization* dataset. Exponential fits are shown as dashed line, with the fitted values between brackets in the legend.

fieldsOfStudy	#Publications		
(Computer Science,)	99230		
(Political Science,)	49987		
(Engineering,)	37469		
(Medicine,)	29223	1st_field	#Publications
(Business,)	29134	Computer Science	121600
(Economics,)	26716	Political Science	58140
(Sociology,)	15869	Engineering	48458
(Geography,)	13211	Business	35955
(Mathematics,)	8192	Medicine	35639
(Art,)	6474	Economics	31842
(Philosophy,)	5767	Sociology	18713
(Environmental Science,)	5708	Geography	15746
(Engineering, Computer Science)	4482	Mathematics	12563
(History,)	3941	Art	7521
(Psychology,)	3702	Environmental Science	6974
(Physics,)	3406	Philosophy	6638
(Mathematics, Computer Science)	2633	Psychology	5115
(Computer Science, Mathematics)	2599	Biology	4944
(Computer Science, Engineering)	2362	History	4587
(Biology,)	2198	Physics	4355
(Biology, Medicine)	2025	Chemistry	2484
(Materials Science,)	2007	Materials Science	2477
(Computer Science, Medicine)	1825	Geology	1393
(Chemistry,)	1394		
(Geology,)	1194		
(Business, Medicine)	1167		
(Business, Computer Science)	1070		

Table D.2: **Number of papers in decreasing order in the dataset that have been categorized with the respective tuple of fields of study (left table) and with the respective first field (right table).** Here we have listed only the fields of study with more than 1000 papers in the dataset. Tuples with more than two fields of study have been reduced considering only the first two.

## D.2 Hierarchical clustering

### D.2.1 Keywords annotation

As described in the main text, the hSBM algorithm produces a hierarchical clustering of both articles and words (i.e. topics). Here we give more details on the procedure we used to manually assign keywords to represent the different document clusters. First, for each cluster at level 4, we manually inspect the most frequent words in the publications' titles. We then look at the most significant topics represented in the articles, as quantified by the normalized mixture proportion [279]. For instance, as we can see depicted in Fig. D.2(b) for the blockchain cluster, or in Fig. D.3(b) for the governance cluster, we can visualize the hierarchical tree of topics, highlighting the top 10 most significant ones. We have also reported in the tree the significance level ( $> 1$ ) with the \* symbol. Moreover, for each of these topic, the top 20 most significant words are printed. We also look at the fields of study represented in each cluster in time, as shown in Fig. D.2(a) for the blockchain cluster, or in Fig. D.3(a) for the governance cluster. Finally, we look at the most important articles in the cluster, as represented by different metrics: overall number of citations, number of citations within the *(de)centralization* dataset, highest knowledge flows towards other clusters, and most central publications according to different network centrality measures (degree, pagerank, betweenness, closeness, katz).

For clarity, we have then chosen one representative keyword for each cluster at the 3<sup>rd</sup> level. In Table D.3, we show the results of the keyword annotation procedure. Starting from the leftmost column, we show each cluster at level 3 with its representative keyword and the number of publications in it, to then show the annotated keywords in each cluster at level 4 (and the number of publications) present in its hierarchical branch in separate rows. Clusters with less than 500 documents have been disregarded.

### Blockchain cluster at level 3




Figure D.2: **Keywords annotation for Blockchain.** (a) Number of papers in time in each field of study for the Blockchain cluster. (b) Hierarchical topic tree, highlighting the top 20 words of the 10 most significant topics in the Blockchain cluster according to the normalized mixture proportion (in color). This information is used to aid the keywords annotation manual procedure.

### Governance cluster at level 3




Figure D.3: **Keywords annotation for Governance.** (a) Number of papers in time in each field of study for the Governance cluster. (b) Hierarchical topic tree, highlighting the top 20 words of the 10 most significant topics in the Governance cluster according to the normalized mixture proportion (in color). This information is used to aid the keywords annotation manual procedure.

	keyword	count	sub-keywords	count
0	Robot swarms	5843	routing, allocation, congestion	909
			robot swarms, supply chain	4120
			planning, game theory	693
1	Cybersecurity	18437	cybersecurity, peer2peer, fault tolearance, attacks	15189
			cloud, security	3248
2	Smart grids	867	smart grids, energy trading, blockchain	867
3	Blockchain	9637	blockchain, cryptocurrency, ethereum, Bitcoin	9634
4	Federated learning	2043	federated learning, deep learning, adversarial networks	1565
5	Optimization	5878	control theory, equilibrium	926
			statistical learning, optimization, detection	3065
			heterogeneity, links, web, health	1887
6	Control theory	19758	control theory, nonlinear dynamics	8634
			navigation, flocking, formation	9661
			discrete-event systems, decision making	1463
7	Electricity	8090	electricity, grids	8090
8	Investments	8394	networks, connectivity, synchronization, routing, topology	3058
			investments, money, risk, algorithm	5336
9	Edge-computing	2970	edge-computing, cloud-computing	2970
10	Telecommunication	18804	cellular networks, communication, radio	12446
			wireless, localization	1262
			routing, protocols, security	5096
11	Wireless technologies	2166	spanish	1251
			wireless, cybersecurity, allocation, decision making	915
12	Governance	40045	governance, fiscal federalism, government, development	19371
			natural resources, education, governance	20674
13	Environment	15386	routing	1278
			renewable energy, wastewater, environment	8348
			portoguese, latin america	5760
14	Social network analysis	10375	supply chain, manufacturing, pricing	3938
			social network analysis, organizations, firms	6437
15	Health	12343	algebra, healthcare, symptoms	6291
			hospitals, HIV, cancer, surgery	6052

Table D.3: **Results of the keywords annotation procedure.** Clusters at 3<sup>rd</sup> level, with associated manually annotated keyword and number of publications, together with the same information for each of the clusters at hierarchical level 4 in its branch of the hierarchy. Single horizontal lines denote clusters at the 2<sup>nd</sup> hierarchical level, while the double horizontal lines denote the three branches at the 1<sup>st</sup> level. Clusters with less than 500 documents have been disregarded.

## D.2.2 Results of the hSBM

In this section we report some additional visualizations of the results of the hSBM algorithm on the *(de)centralization* dataset. In Fig. D.4 we show the number of papers in each cluster at hierarchical level 4, together with the hierarchical tree and the manually annotated keywords. In Fig. D.5 we show the bipartite network of documents and words, together with the results of the hSBM algorithm and the manual annotation procedure. This figure shows how the various clusters are represented in terms of topic frequency. Nodes on the left, indeed, represent documents, while those on the right title words. Notice that, to properly represent topic frequencies, multiple instances of the same word are considered, one for each document the word is in. Links are colored based on the doc cluster they start from. On top of the bipartite network between docs and words, we show a tree where squares represent clusters/topics at various hierarchical levels as a result of the hSBM, starting from the common root to level 3.





Figure D.4: **Results of hSBM at the 4<sup>th</sup> hierarchical level.** Heatmap of the number of papers in each cluster at level 4, together with the manually annotated keywords and the hierarchical tree resulting from the clustering algorithm.



- T0 (control, optimal, access, decentralised, cooperative, method, robust, peer, controller, multiple)
- T1 (based, management, study, service, case, strategy, evaluation, review, community, towards)
- T2 (economic, urban, low, heterogeneous, web, organization, environmental, city, uncertain, business)
- T3 (technology, flow, descentralización, em, décentralisation, technologies, na, descentralização, saúde, brasil)
- T4 (power, energy, centralized, information, supply, state, chain, coordination, planning, impact)
- T5 (performance, process, centralization, treatment, use, industrial, driven, cancer, project, centralizers)
- T6 (decentralization, local, development, health, governance, public, social, fiscal, government, china)
- T7 (data, blockchain, sensor, mobile, resource, framework, smart, architecture, grid, allocation)
- T8 (network, distributed, networks, multi, analysis, agent, algorithm, learning, dynamic, policy)
- T9 (decentralized, system, using, systems, approach, model, wireless, design, time, large)
- T10 (dual, exchange, loop, trading, closed, sensitive, reverse, remanufacturing, transactive, commodity)
- T11 (new, demand, key, evolution, autonomy, action, identity, transformation, authority, motion)

Figure D.5: **Visualization of the results of the hSBM algorithm.** Bipartite network of documents and words, where nodes and links are colored based on the document cluster (at hierarchical level 3) they start from. Multiple instances of the same word are considered, one for each document the word is in. Word nodes within a topic at level 3 are ordered as the document clusters for visualization purposes. On top of the bipartite network, the hierarchy of topics and document clusters are represented on the word and document side respectively, shown as a network of square gray nodes. On the bottom of the figure, the legend with the 10 most frequent words of each topic is displayed.

### D.3 The importance of governance and blockchain in the history of (de)centralization

In this section we report some additional and complementary plots to show the importance of the Blockchain and Governance clusters in the academic literature on *(de)centralization*. In Fig. D.6(a) we show the number of papers in time for three groups of documents: Blockchain, Governance and the whole dataset. Moreover, in Fig. D.6(b) we plot the rank of each cluster in terms of yearly number of papers, highlighting Blockchain and Governance. Both figures clearly show how Governance has been the most productive cluster in the literature for a long time, only to be replaced by Blockchain in recent years.

In Fig. D.7 we look at knowledge flows, plotting their value instead of their rank (as done in Fig. 3 in the main text). We plot the time evolution of  $K_{a \rightarrow \bullet}(Y)$ , where  $a$  is Governance or Blockchain, in comparison to  $K_{\bullet \rightarrow \bullet}(Y)$  in Fig. D.7(a). This shows the average influence that the cluster in a certain year has towards the other clusters in the future. We similarly compare the average influence of all other clusters on the future of these two clusters looking at  $K_{\bullet \rightarrow a}(Y)$  in Fig. D.7(b). We can see from these figures how Governance has been increasingly important in influencing other clusters until the 1980s, while since after the 1990s it has had a lower knowledge flow than the average among all clusters, despite being the first cluster in terms of number of papers in all this time. The case of Blockchain is opposite: after 2013 it starts to have a much higher influence towards the other clusters compared to the average one over all clusters, while the other clusters after 2004 have had an average knowledge flow towards it.




Figure D.6: **Blockchain and Governance are the most productive clusters.** (a) Number of papers in time for the Blockchain cluster, the Governance cluster and the full *(de)centralization* dataset. (b) Rank by number of papers in time for different clusters, highlighting Blockchain and Governance. Both plots show the central role of these two clusters in the literature on *(de)centralization*, and how they have exchanged roles, with Blockchain becoming the most productive field in recent years.




Figure D.7: **Influence of Governance and Blockchain from and on other clusters in time.** Average knowledge flow towards other clusters in the future in time from a cluster (a) or to a cluster (b), the latter being either Blockchain (in blue) or Governance (in red), while in gray the average over the *(de)centralization* dataset is plotted.

### D.3.1 Influence of governance from and on other clusters

Here we report an additional plot looking at the role of Governance with respect to the other clusters. In particular, we replicate Fig. 4 in the main text for Governance instead of Blockchain. Given the results on the predominant role of Governance as an influence on other clusters before the year 2000, we look at the following three periods of time: from 1950 to 1980, from 1981 to 1990 and from 1991 to 2000, highlighting different interactions with other clusters both in terms of source and destination of knowledge flows. For instance, the Blockchain cluster is the second most influenced cluster by Governance in the 1990s, highlighting its role in the foundations of the field. However, the rankings are generally more stable than what observed in Fig. 4 in the main text for the case of the Blockchain cluster.




Figure D.8: **The influence from/on Governance with regards to the rest of the (de)centralization literature.** (a) Change of the ranking of the clusters most influenced by the Governance literature between its early period (1950-1980), its middle period (1981-1990), and its late period (1991-2000), calculated using the average knowledge flows  $K_{a \rightarrow b}(T)$ , where  $T$  is the selected period, and  $a$  is fixed to Governance. (b) Change of the ranking of the most influential clusters on the Governance literature between its early period (1950-1980), its middle period (1981-1990), and its late period (1991-2000), calculated using the average knowledge flows  $K_{a \rightarrow b}(T)$ , where  $T$  is the selected period, and  $b$  is fixed to Governance. In both cases, if  $K_{a \rightarrow b}(T) = 0$ , we print a circle in the corresponding gray node and use a lighter color in the respective link. Moreover, we print a star when  $0.01 < K_{a \rightarrow b}(T) \leq 0.1$ , and two stars when  $K_{a \rightarrow b}(T) > 0.1$ .

# Chapter E

## Appendix to chapter 7

### E.1 Additional methods

**Evaluation of coefficients of the trend line in Figure 7.2(a).** The coefficients  $a = 1.06$  and  $b = 0.70$  of the trend line  $y = x^a 10^{-b}$  in Figure 7.2(a) are in good agreement with the empirical data,  $R^2 = 0.969$ , and evaluated as follows. First, the equation is transformed to  $Y = aX - b$ , where  $Y = \log_{10} y$  and  $X = \log_{10} x$ . The linear equation fitted against real data and coefficients  $a$  and  $b$  computed by minimizing the sum of squares.

**Statistical analysis.** We compare the median of two paired distributions using the two-sided Wilcoxon test [309]. It is a non-parametric statistical test and verifies the null hypothesis that two paired samples come from distributions with the same median. If distributions are not paired, we use the Mann-Whitney-U test to assess statistical differences of the medians of two distributions [311]. We compare two distributions using the Kolmogorov-Smirnov test [312] on two samples. It tests the null hypothesis that 2 independent samples are drawn from the same continuous distribution. We evaluate the correlation between two sets of values using the Spearman rank-order correlation coefficient [313]. It is a correlation coefficient that does not assume normally distributed values and varies between -1 and 1: with -1 implying a negative correlation, 0 no correlation, and 1 a positive correlation.

## E.2 General statistics of the 40 DWMs under consideration

Name	Transactions with a DWM		U2U transactions	
	Users (sent; received; total)	Trading volume in millions (sent; received; total)	Users (sent; received; total)	Trading volume in millions (total)
Abraxas	(95,642; 21,500; 111,003)	(21.85; 27.23; 49.09)	(28,588; 25,546; 44,151)	61.92
Agora	(462,106; 119,221; 537,983)	(141.3; 132.8; 274.1)	(168,248; 151,699; 252,984)	558.0
AlphaBay	(1,658,059; 334,154; 1,898,850)	(537.1; 568.6; 1,106)	(524,783; 422,881; 776,183)	1581
Apollon	(68,373; 13,954; 79,307)	(12.90; 16.59; 29.50)	(19,468; 17,290; 29,900)	49.38
Basetools	(119,114; 347; 119,461)	(4.712; 6.727; 11.44)	(32,191; 34,169; 50,939)	63.23
Benumb Shop	(27,229; 343; 27,556)	(3.929; 5.027; 8.956)	(5,499; 5,654; 8,985)	21.73
BitBazaar	(20,805; 150; 20,931)	(2.681; 4.425; 7.106)	(6,939; 6,569; 10,126)	14.13
Black Bank	(52,783; 15,147; 64,131)	(11.41; 11.78; 23.19)	(15,843; 13,486; 24,291)	31.11
Blue Sky	(16,002; 10,140; 22,616)	(3.225; 3.786; 7.011)	(9,763; 6,149; 12,108)	10.86
Buybest	(334,741; 3,004; 337,556)	(24.45; 7.490; 31.94)	(57,001; 59,131; 99,390)	132.4
Bypass Shop	(861,716; 8,118; 869,593)	(65.66; 54.36; 120.0)	(176,905; 174,151; 288,745)	804.3
DarkMarket	(176,141; 13,554; 183,010)	(27.25; 36.62; 63.87)	(72,923; 67,621; 105,416)	166.7
Dream	(466,511; 45,399; 507,837)	(57.64; 72.95; 130.6)	(109,871; 70,706; 154,873)	287.3
Empire	(405,202; 9,690; 413,858)	(64.38; 56.84; 121.2)	(63,431; 56,415; 103,886)	287.4
Evolution	(216,604; 34,512; 240,713)	(47.58; 50.15; 97.73)	(77,331; 72,711; 115,496)	236.7
FEshop	(1,134,456; 5,858; 1,140,275)	(64.67; 48.83; 113.5)	(244,318; 261,489; 420,040)	834.4
Flugsvamp 2.0	(104,385; 21,201; 119,893)	(23.01; 38.20; 61.22)	(29,215; 23,079; 41,047)	144.2
Flugsvamp 3.0	(217,083; 20,773; 234,563)	(39.34; 52.78; 92.12)	(52,075; 49,881; 81,527)	473.8
FuLLzShOp	(21,716; 9; 21,726)	(3.937; 4.510; 8.447)	(4,147; 4,496; 7,209)	10.07
Hansa	(330,565; 73,202; 358,120)	(60.64; 55.91; 116.6)	(153,567; 127,514; 209,717)	76.60
Hydra	(4,031,013; 666,075; 4,584,339)	(1,868; 1,810; 3,678)	(2,447,548; 2,099,320; 3,124,366)	20,840
Joker's Stash	(806,089; 1,090; 807,140)	(153.0; 49.95; 203.0)	(154,872; 156,689; 260,832)	926.2
LuxSocks.ru	(326,159; 186; 326,340)	(8.123; 5.573; 13.70)	(59,638; 66,011; 97,705)	175.8
Matanga	(57,354; 633; 57,963)	(5.882; 7.775; 13.66)	(10,637; 10,328; 17,632)	96.35
Middle Earth	(38,017; 9,206; 45,629)	(8.361; 9.151; 17.51)	(8,091; 7,603; 12,990)	18.68
MrGreen.ws	(44,918; 176; 45,094)	(8.244; 6.176; 14.42)	(6,298; 5,912; 10,501)	14.44
Nightmare	(37,844; 3,524; 40,894)	(5.697; 7.371; 13.07)	(8,830; 6,277; 12,905)	25.83

Table E.1: **General statistics of DWMs, part 1.** Some DWMs are presented here, the others are available in Table E.2. The terms “sent” and “received” always refer to transactions made by users. The trading volume indicates millions of dollars. The amount of dollars sent and received by users through U2U transactions is equivalent to the total.

Name	Interactions with DWM		U2U interactions	
	Users (sent; received; total)	Trading volume in millions (sent; received; total)	Users (sent; received; total)	Trading volume in millions (total)
Nucleus	(205,043; 53,571; 247,884)	(56.59; 61.68; 118.3)	(62,577; 52,829; 93,279)	156.9
Pandora	(35,667; 8,723; 41,718)	(8.401; 8.561; 16.96)	(11,119; 8,964; 15,944)	26.37
Russian Anonymous	(740,625; 36,161; 769,228)	(80.24; 95.94; 176.2)	(363,773; 331,811; 493,766)	1866
San-Wells	(51,795; 2,858; 54,633)	(6.335; 5.755; 12.09)	(8,227; 7,841; 13,679)	15.36
Sheep	(38,068; 7,634; 42,673)	(10.81; 11.47; 22.29)	(12,007; 10,288; 182,90)	49.87
Silk Road	(382,534; 72,344; 429,284)	(130.2; 149.7; 279.9)	(163,376; 157,113; 243,441)	671.4
Silk Road 2	(222,666; 47,528; 254,830)	(66.83; 71.92; 138.7)	(73,116; 66,019; 111,387)	259.8
Silk Road 3.1	(59,894; 15,413; 70,078)	(9.054; 13.49; 22.54)	(22,160; 18,570; 32,491)	21.80
TradeRoute	(103,517; 14,080; 112,634)	(16.97; 17.04; 34.01)	(27,901; 22,287; 41,869)	67.72
Unicc	(2,004,236; 559; 2,004,789)	(147.8; 84.61; 232.4)	(473,969; 490,794; 780,282)	1,673
Valhalla	(82,507; 8,218; 89,214)	(8.933; 9.811; 18.74)	(25,755; 32,687; 45,297)	51.49
Wall Street	(334,871; 25,352; 347,842)	(48.15; 53.16; 101.3)	(148,262; 127,370; 203,176)	163.5
xDedic	(27,956; 885; 28,736)	(3.552; 3.838; 7.389)	(4,785; 4,767; 7,685)	12.70

Table E.2: **General statistics of DWMs, part 2.** Some DWMs are presented here, the others are available in Table E.1. The terms “sent” and “received” always refer to transactions made by users. The trading volume indicates millions of dollars. The amount of dollars sent and received by users through U2U transactions is equivalent to the total.




Figure E.1: **Lifetime of DWMs in our dataset.** Time interval between the first and last transaction of each DWM. A total of 17 DWMs participated in at least one transactions in either 2020 or 2021, while 23 closed before 2020.




Figure E.2: **Key distributions of the full network.** Probability distribution function (pdf) about the number of transactions of each pair of entities (a)-(d)-(g), their trading volume (b)-(g)-(h), and their lifetime computed as time difference between their last and first transaction (c)-(f)-(i).

## E.3 Detection of stable pairs in temporal and directed networks

Here, we summarize the methodology of detecting the backbone of stable pairs in temporal and undirected networks as introduced in [1], and show how it can be easily adapted to tackle the analysis of directed temporal networks. The methodology follows three sequential steps: (i) determine the interval partition, (ii) estimate models' parameters, over successive intervals, and (iii) run a statistical filter, which removes all pairs explained by the null hypothesis and retain stable pairs. The analysed temporal network, either directed or undirected, of  $N$  nodes evolves in an observation window composed of  $T \gg 1$  time steps, labeled as  $t = 1, \dots, T$ . At each time step  $t$ , entities interact among themselves and form a time-varying network of interactions, described by a binary adjacency matrix that varies in time  $A(t)$ .

### E.3.1 Temporal and undirected networks

**Interval partition.** The overall observation window is divided in successive and disjoint intervals using an auxiliary method, namely, the Bayesian Block method [353]. It takes as input the total number of temporal pairs created in the entire network at time  $t$

$$\Omega^{\text{ts}}(t) = \sum_{i,j=1;i \leq j}^N A_{ij}^{\text{ts}}(t), \quad (\text{E.1})$$

where the superscript “ts” indicates that these variables are estimated from the time series and  $A_{ij}^{\text{ts}}(t)$  is the  $ij$ th entry of the estimated adjacency matrix at time  $t$ . The Bayesian Block method returns the interval partition, which divides the overall time window  $T$  into  $I$  disjoint intervals indexed by  $\Delta = 1, \dots, I$ , that contain a uniform total number of connections. From the knowledge of the interval partition, the length,  $\tau(\Delta)$ , of the generic  $\Delta$ th interval is obtained with the following closure relation:  $\sum_{\Delta=1}^I \tau(\Delta) = T$ .

**Parameter estimation.** According to the null hypothesis, pair of entities  $i$  and  $j$  are expected to interact proportional to their individual activities at time  $t$ . That is, the probability that entities  $i$  and  $j$  interact at time  $t$  is a binomial random variable defined as

$$p_{ij}(t) \equiv a_i(t)a_j(t), \quad (\text{E.2})$$

where  $a_i(t)$  and  $a_j(t)$  are piece-wise constant activities, which represent the propensity of creating interactions at time  $t$ . The estimation of piece-wise constant activities is carried out analysing each of the  $I$  intervals separately. The activity of entity  $i$  at time  $t \in [t_{\text{in}(\Delta)}, t_{\text{in}(\Delta)} + \tau(\Delta) - 1]$  is computed through the following frequency count:

$$a_i(t) = \frac{s_i^{\text{ts}}(\Delta)}{\sqrt{2W^{\text{ts}}(\Delta)\tau(\Delta)}}, \quad (\text{E.3})$$

where  $s_i^{\text{ts}}(\Delta)$  and  $W^{\text{ts}}(\Delta) \gg 1$  are the total number of pairs generated by entity  $i$  in the  $\Delta$ th and the total number of temporal pairs generated in the network in the  $\Delta$ th interval, respectively. These variables are computed from the adjacency matrix  $A^{\text{ts}}(t)$ , as,  $s_i^{\text{ts}}(\Delta) = \sum_{j=1}^N \sum_{t=t_{\text{in}}(\Delta)}^{t_{\text{in}}(\Delta)+\tau(\Delta)-1} A_{ij}^{\text{ts}}(t)$ , and  $W^{\text{ts}}(\Delta) = \frac{1}{2} \sum_{i=1}^N s_i^{\text{ts}}(\Delta)$ . Once the activities are estimated according with Eq. (E.3), the probability in Eq. (E.2) can be calculated.

**Statistical filter.** The statistical filter compares expected number of connections between entity  $i$  and entity  $j$ ,  $E[\bar{w}_{ij}]$ , with observations from the time series,  $\bar{w}_{ij}^{\text{ts}} = \sum_{t=1}^T A_{ij}^{\text{ts}}(t)$ . The expected number of connections between entities  $i$  and  $j$  in the overall time window  $T$  is determined by the sum of the binomial random variables given in Eq. (E.2)

$$E[\bar{w}_{ij}] = \sum_{t=1}^T p_{ij}(t) = \sum_{\Delta=1}^I \frac{s_i^{\text{ts}}(\Delta) s_j^{\text{ts}}(\Delta)}{2W^{\text{ts}}(\Delta)}, \quad (\text{E.4})$$

where we have used the estimation of activity in Eq. (E.3) and summed over all intervals. Although the sum of non-identical binomial random variables in Eq. (E.4) is a Poisson binomial distribution, the Poisson distribution is an appropriate approximation for long time series. The probability that the observed weight,  $\bar{w}_{ij}^{\text{ts}}$ , could be explained by the relative expected weight,  $E[\bar{w}_{ij}]$  in Eq. (E.4), is computed according to the cumulative function of the Poisson distribution

$$\alpha_{ij} \equiv 1 - \sum_{x=0}^{\bar{w}_{ij}^{\text{ts}}-1} P(x; E[\bar{w}_{ij}]), \quad (\text{E.5})$$

where  $P(x; E[\bar{w}_{ij}])$  indicates the Poisson distribution with random variable  $x$  and expected value  $E[\bar{w}_{ij}]$ . Equation (E.5) represents the p-value  $\alpha_{ij}$ : when the p-value is below a pre-defined threshold, the pair  $ij$  is significant and included in the backbone network. The same statistical test is repeated for all pairs of entities  $ij$  observed at least once in the overall temporal evolution.

### E.3.2 Temporal and directed networks

With little modifications, the above methodology can be used to filter temporal and directed networks.

**Interval partition.** The interval partition is obtained by using the Bayesian Block method as above. The total number of temporal pairs created in the entire network at time  $t$  is

$$\Omega^{\text{ts}}(t) = \sum_{i,j=1}^N A_{ij}^{\text{ts}}(t), \quad (\text{E.6})$$

where not pairs are directed, while in Eq. (E.1) undirected, thereby explaining the different ranges in the summations.

**Parameter estimation.** In directed networks, the probability that entity  $i$  contacts at random entity  $j$  at time  $t$  is defined as

$$p_{i \rightarrow j}(t) \equiv a_i(t)b_j(t). \quad (\text{E.7})$$

where  $a_i(t)$  is the activity of entity  $i$  at time  $t$  and  $b_j(t)$  the attractiveness of entity  $j$  at time  $t$ . The activity was already defined in Eq. (E.2), while the attractiveness represent the propensity of receiving connections at time  $t$ . If  $a_i(t) = b_i(t) \forall i, t$  (for all entities in the network and at all time), Eq. (E.7) becomes equivalent to Eq. (E.2). However, care should be placed in their interpretation, whereby Eq. (E.7) generates a directed pair from entity  $i$  to entity  $j$ , while Eq. (E.2) can only lead to an undirected pair.

In the generic  $\Delta$ th interval, defining the time window  $t \in [t_{\text{in}(\Delta)}, t_{\text{in}(\Delta)} + \tau(\Delta) - 1]$ , piecewise constant activities and attractivenesses are estimated directly from the time series, similarly to what done in the undirected case in Eq. (E.3)

$$a_i(t) = \frac{s_{\text{out},i}^{\text{ts}}(\Delta)}{\sqrt{W^{\text{ts}}(\Delta)\tau(\Delta)}} \quad b_i(t) = \frac{s_{\text{in},i}^{\text{ts}}(\Delta)}{\sqrt{W^{\text{ts}}(\Delta)\tau(\Delta)}}, \quad (\text{E.8})$$

where  $s_{\text{out},i}^{\text{ts}}(\Delta)$ ,  $s_{\text{in},i}^{\text{ts}}(\Delta)$ , and  $W^{\text{ts}}(\Delta) \gg 1$ , are the total incoming strength of entity  $i$  in the  $\Delta$ th interval, outgoing strength of entity  $i$  in the  $\Delta$ th interval, and the total number of directed, temporal pairs generated in the network in the  $\Delta$ th interval, respectively. These variables are computed from the adjacency matrix  $A^{\text{ts}}(t)$ , that is,  $s_{\text{out},i}^{\text{ts}}(\Delta) = \sum_{j=1}^N \sum_{t=t_{\text{in}(\Delta)}}^{t_{\text{in}(\Delta)}+\tau(\Delta)-1} A_{ij}^{\text{ts}}(t)$ ,  $s_{\text{in},i}^{\text{ts}}(\Delta) = \sum_{i=1}^N \sum_{t=t_{\text{in}(\Delta)}}^{t_{\text{in}(\Delta)}+\tau(\Delta)-1} A_{ij}^{\text{ts}}(t)$ , and  $W^{\text{ts}}(\Delta) = \sum_{i=1}^N s_{\text{out},i}^{\text{ts}}(\Delta)$ . Once the activity and attractiveness are estimated according with Eq. (E.8), the probability in Eq. (E.7) can be evaluated.

**Statistical filter.** Similar to Eq. (E.4), the expected number of pairs from entity  $i$  to entity  $j$  is computed by summing the probability in Eq. (E.7) for all time instants  $t$

$$\text{E}[\bar{w}_{i \rightarrow j}] = \sum_{t=1}^T p_{i \rightarrow j}(t) = \sum_{\Delta=1}^I \frac{s_{\text{out},i}^{\text{ts}}(\Delta)s_{\text{in},j}^{\text{ts}}(\Delta)}{W^{\text{ts}}(\Delta)}. \quad (\text{E.9})$$

The probability that the observed weight,  $\bar{w}_{i \rightarrow j}^{\text{ts}}$ , is explained by the expected weight,  $\text{E}[\bar{w}_{i \rightarrow j}]$  in Eq. (E.9), is computed according to the cumulative function of the Poisson distribution

$$\alpha_{i \rightarrow j} \equiv 1 - \sum_{x=0}^{\bar{w}_{i \rightarrow j}^{\text{ts}}-1} P(x; \text{E}[\bar{w}_{i \rightarrow j}]). \quad (\text{E.10})$$

Equation (E.10) represents the p-value  $\alpha_{i \rightarrow j}$ , which is used to assess whether the directed pair  $i \rightarrow j$  is significant. The same statistical test has to be repeated for directed pairs observed at least once in the overall temporal evolution. For undirected networks, Eq. (E.10) is equivalent to Eq. (E.5).

## E.4 Additional simulations




Figure E.3: **Statistics of U2U pairs.** (a) Number of stable and non-stable U2U pairs with a given number of transactions. (b) Number of stable and non-stable U2U pairs with a given trading volume. Vertical lines represent median values of the respective distributions.




Figure E.4: **Evaluation of the time users spent on a DWM.** It extends Figure 7.4(inset) in the main text by considering each individual DWM. Statistical tests are carried using the two-sided Kolmogorov-Smirnov test and results are available in Table E.3. Vertical lines represent median values of the respective distributions.




Figure E.5: **Evaluation of the total trading volume users exchange with a DWM.** It extends Figure 7.4 in the main text by considering each individual DWM. Statistical tests are carried using the two-sided Kolmogorov-Smirnov test and results are available in Table E.3. Vertical lines represent median values of the respective distributions.

Name	Time spent on a DWM	Trading volume exchanged with a DWM
	Users with stable U2U pairs vs other users (KS; p-value)	Users with stable U2U pairs vs other users (KS; p-value)
Abraxas	(0.529; 0.0001)	(0.355; 0.0001)
Agora	(0.583; 0.0001)	(0.351; 0.0001)
AlphaBay	(0.561; 0.0001)	(0.327; 0.0001)
Apollon	(0.540; 0.0001)	(0.382; 0.0001)
Basetools	(0.504; 0.0001)	(0.394; 0.0001)
Benumb	(0.519; 0.0001)	(0.286; 0.0001)
BitBazaar	(0.482; 0.0001)	(0.394; 0.0001)
Black Bank	(0.409; 0.0001)	(0.290; 0.0001)
Blue Sky	(0.535; 0.0001)	(0.346; 0.0001)
Buybest	(0.534; 0.0001)	(0.338; 0.0001)
Bypass	(0.584; 0.0001)	(0.426; 0.0001)
DarkMarket	(0.635; 0.0001)	(0.459; 0.0001)
Dream	(0.475; 0.0001)	(0.450; 0.0001)
Empire	(0.491; 0.0001)	(0.395; 0.0001)
Evolution	(0.557; 0.0001)	(0.305; 0.0001)
FEshop	(0.639; 0.0001)	(0.447; 0.0001)
Flugsvamp 2.0	(0.551; 0.0001)	(0.423; 0.0001)
Flugsvamp 3.0	(0.604; 0.0001)	(0.500; 0.0001)
FuLLzShOp	(0.513; 0.0001)	(0.222; 0.0001)
Hansa	(0.554; 0.0001)	(0.359; 0.0001)
Hydra	(0.536; 0.0001)	(0.327; 0.0001)
Joker's Stash	(0.647; 0.0001)	(0.389; 0.0001)
LuxSocks.ru	(0.676; 0.0001)	(0.390; 0.0001)
Matanga	(0.544; 0.0001)	(0.405; 0.0001)
Middle Earth	(0.473; 0.0001)	(0.317; 0.0001)
MrGreen.ws	(0.537; 0.0001)	(0.280; 0.0001)
Nightmare	(0.401; 0.0001)	(0.400; 0.0001)
Nucleus	(0.559; 0.0001)	(0.380; 0.0001)
Pandora	(0.505; 0.0001)	(0.313; 0.0001)
Russian Anonymous	(0.682; 0.0001)	(0.533; 0.0001)
San-Wells	(0.558; 0.0001)	(0.302; 0.0001)
Sheep	(0.472; 0.0001)	(0.344; 0.0001)
Silk Road 2	(0.539; 0.0001)	(0.365; 0.0001)
Silk Road 3.1	(0.542; 0.0001)	(0.415; 0.0001)
Silk Road	(0.589; 0.0001)	(0.341; 0.0001)
TradeRoute	(0.496; 0.0001)	(0.386; 0.0001)
Unicc	(0.708; 0.0001)	(0.537; 0.0001)
Valhalla	(0.448; 0.0001)	(0.440; 0.0001)
Wall Street	(0.562; 0.0001)	(0.405; 0.0001)
xDedic	(0.546; 0.0001)	(0.172; 0.0001)

Table E.3: **Statistical tests.** The two-sided Kolmogorov-Smirnov test is used to perform the statistical test. All p-values are less than 0.0001, which is indicated with 0.0001.




Figure E.6: **DWMs where users meet.** Scatter plot of the number of pairs of users that meet inside each of the 40 DWMs considered versus the total volume sent to the DWM.




Figure E.7: **Trading volume of stable U2U pairs during COVID-19.** 28-days moving average of trading volume between users who met inside a DWM (a) and outside any DWMs (b). (c) Yearly growth relative to the same day of 2019. Vertical lines represent the dates of Wuhan (Jan 23, 2020) and Italy (March 3, 2020) lockdowns.




Figure E.8: **Formation of new stable U2U pairs during COVID-19.** 28-days moving average of new stable U2U pairs started between U2U pairs who met inside a DWM (a) and outside any DWMs (b). (c) Yearly growth relative to the same day of 2019. Vertical lines represent the dates of Wuhan and Italy lockdowns.

# Chapter F

## Appendix to chapter 8

### F.1 OpenSea market mechanisms

For a majority of NFT collectibles, the minting happens as follows. The creators offer the possibility for anyone with a wallet to generate a new NFT for a fixed price, whose attributes will be randomly selected, even though each attribute can only be given to a specific amount of NFTs. Once every NFT has been minted by the community, they are made available to their buyers, who can sell them on a marketplace afterwards.

Before releasing their collection, creators also set how much royalty they want to get from each secondary sale related to their NFTs. As such, every time a new sale happens, the royalty is deducted from the share the seller gets, as well at 2.5% of the total price that OpenSea gets from every sale taking place on their platform.

The following table details, for a few collections, the initial price at which the NFTs could be minted (gas fee, i.e., the fees required to conduct a transaction on the Ethereum blockchain, not included). Note that these transactions are not considered as sales per se by OpenSea's official API.

<b>Collection</b>	<b>Minting Price</b>
CryptoPunks	Free
Bored Ape Yacht Club	0.08 ETH
World of Women	0.07 ETH
CryptoTrunks	0.5 ETH
CryptoCorgi	First corgi to be claimed at 0.001 ETH, last one at 1.001 ETH
Sewer Rat Social Club	0.05 ETH
Rabbit College Club	0.02 ETH
Cute Pig Club	0.03 ETH
Ape Gang	Free

## F.2 Generative art mechanisms

As previously mentioned, NFT collectibles are usually generated using an algorithmic procedure, which can lead to thousands of unique tokens created with the same set of instructions [354]. However, the inner workings of the algorithms have not been shared by the creators, and can greatly differ between collections. It is therefore impossible to assess whether the rarity curves for the collections displayed in Section “Quantifying Rarity” share similarities because their algorithms follow similar steps. In the case of the CryptoPunks, members of the community have been attempting to reverse-engineer the algorithm used by Larva Labs to generate the original Punks [347], or even to replicate it [355, 356]. However, the creators never released any information on the matter, as well as any other NFT collectibles creator.

## F.3 List of collections

Collection Names				
0N1 Force	0xVampire Project	24px	8 BIT UNIVERSE	Absurd Arboretum
Adam Bomb Squad	AfroDroids By Owo	AI Cabones	AlphaBetty Doodles	AmeegosOfficialNFT
Angels of Aether	Angry Boars	AnimalWorldWar	Animathereum	Animetas
Ape Gang	Ape Harbour Yachts	ApesOfSpace	Approving Corgis	Arabian Camels
ArcadeNFT	Art Stars Club Official	Astro Frens	Astrohedz	Avarik Saga Universe
Avastars	Axolittles	BASTARD GAN PUNKS V2	BLU Blox	BULLSEUM
BYOPills	Baby Combat Bots G1	Bad Bunnies NFT	Bad Kids Alley Official	Badass Bulls
Barn Owls	Barn Owls Dino Palz	Based Fish Mafia	Bear Market Bears	Bears Deluxe
BearsOnTheBlock	Beatnik Tiki Tribe	Bit Wine	BlankFace	Blob Mob
BlockchainBikers	Bones & Bananas	Bones Club Heritage	Bonsai by ZENFT	Bored Ape Kennel Club
Bored Ape Yacht Club	Bored Mummy Baby Waking Up	Bored Mummy Waking Up	Boring Bananas Co.	Boss Beauties
BroadcastersNFT	BullsOnTheBlock	Bunker Beasts	Buzzed Bear Hideout	CHIBI DINOS
COVIDPunks!	CanineCartel	Cartlads	Catctus Collectibles	Catshit Crazy
Chads NFT	ChainFaces	Chibi Apes	Chihuahua Gang	Chill Frogs NFT
Chiptos	Chubbies	Ciphersquares Official	Citizens of Bulliever Island	Claylings
CleverGirls NFT	Cool Cats NFT	Crazy Crows Chess Club	Crazy Dragon Corps	Crazy Lizard Army
CrazySkullzNFT	Criminal Donkeys	Crumbys Bakery	CrypToadz by GREMPLIN	Cryptiniis
Crypto Cannabis Club	Crypto Corgis	Crypto Duckies	Crypto Ghosts NFT	Crypto Hobos
Crypto Hodlers NFT	Crypto Squatches	Crypto Tuners	Crypto.Chicks	CryptoFighters
CryptoFimney	CryptoMutts	CryptoPunks	CryptoSkulls	CryptoTrunks
Cunning Foxes	Cupcarts Official	Cute Pig Club	CyberKongz	CyberKongz VX
CyberPunkA12	Cybergirl Fashion	Cypher City	Dapper Dinos Karma Collective	Dapper Dinos NFT
Dapper Space Collective	Dead Devil Society	DeadFellaz	DeadHeads	Deadbears Official
Deez Nuts (Official Nuts)	Degen Gang	Degenz	Delisted Tiny Punks	Derpy Birbs
Devious Demon Dudes	Dizzy Dragons	Doge Pound Puppies	DogePirates	Dogs Unchained
Dope Shibas	Dreamloops	DystoPunks	Encryptas	Epic Eagles
Ether Cards Founder	EtherGals	Ethereans Official	Etheremura	Evil Teddy Bear Club
FLUF World	FUD Monsters	FVCK_CRYSTAL//	FameladySquas	Fang Gang
Garners Marketverse Patrons	Fast Food Frens Collection	Fast Food Punks	Fatales	Flowtys
Floyds World	Forgotten Runes Wizards Cult	FoxyFam	Frogs In Disguise	FusionApes
Fxck Face	GLICPIXXVER002 - GRAND COLLECTION	GOATz	GRAYCRAFT2	GRILLZ GANG
Galactic Secret Agency	GalacticApes	Galaxy Fight Club	Galaxy-Eggs	GameOfBlocks
Gator World NFT	Gauntlets	Genesis Block Art	Glue Factory Show	Goblin Goons
Good Guys NFT	Goons of Balatroom	Gorilla Nemesis	Great Ape Society	Guardians of the Metaverse
Gutter Cat Gang	Gutter Rats	HDPunks	HODL GANG	Hammys
HappyLand Gummy Bears Official	HashGuise Gen One	Hashmasks	HatchDracoNFT	Heroes of Evermore
Hewer Clan	HodlHeads	Holy Cows	HypeHippo.io	IMMORTALZ - Ambarly Assassins
Incognito	Kamagang	Keplers Civil Society	KidPunks	Knights of Degen - Knights!
Koala Intelligence Agency	Koin Games Dev Squad	Kokeshi World	Krazy Koalas NFT	Lamb Duhs
Lazy Lions	Lazy Lions Bungalows	Lobby Lobsters	Lockdown Lemmings	Lonely Planet Space Observatory
Long Neckie Fellas	Long Neckie Ladies	Loopy Donuty	Loot (for Adventurers)	Lost Souls Sanctuary

Collection Names				
Lostboy NFT	Lucha Libre Knockout	Lucky Maneki	Lucky Sloths NFT	Lumps World
Lysergic Labs Shroomz	MOONDOGS ODYSSEY	Mad Banana Union	Mad Cat Militia	MaestroPups
Magic Mushroom Clubhouse	Mandelbrot Set Collection	Maneki Gang	MarsCatsVoyage	Meebits
Mighty Manateez	Mini Monkey Mafia	Minimints	MissCryptoClub	MjiBots
Monas	MonkePunks	Monkeybrix	Monster Blocks - Official	Monster Rehab 1.0
Mutant Ape Yacht Club	MutantKongz	Muttinks	My Fucking Pickle	NFT Siblings
NFTBOY: Bored Ape Racers	NOOBS NFT	Naughty Tigers Costume Club	Neon Junkies	Nice Drips
Nifty League DEGENs	Niftyriots	Non-Fungible Heroes	Notorious Frogs	ORCZ!
OctoHedz	Oddball Club (Official)	Official DogeX	Omnimorphs	OnChainMonkey
Osiris Cosmic Kids	PEACEFUL GROUPIES	PORK1984	POW NFT	PPPandas
Paladin Pandas	Panda Dynasty	Panda Golf Squad	Party Penguins	Penguin Fight Club
PinapplesDayOut	Pirate Treasure Booty Club	PixaWizards	Platy Punks - Official	PogPunks NFT
Polar Pals Bobsledding	Posh Pandas	Potato Power Club	Primate Social Society	Procedural Space
Pudgy Penguins	PunkBabies	PunkCats	PunkScapes	Purnelopes Country Club
PyMons	Qubits On The Ice	RUUMZ	Rabbit College Club	Raccoon Mafia
Raccoons Club	RagingRhinos	Re-Genz	Ready Player Cat NFT	Reb3l Bots
Reckless Whales	RichKidsOfficial	Rickstro Frens	Rivermen	Roaring Leaders
Robotos Official	Rogue Society Bots	Royal Ceramic Club	Royal Society Chips	Royal Society of Players
Rumble Kong League	SLOTHz	STRAWBERRY.WTF	SVINS	Sad Frogs District
Sad Girls Bar	SamuraiDoge	Sappy Seals	Satoshibles	Savage Droids
Save the Martians	ScoopDog Squad	Secret Society of Whales	Sewer Rat Social Club	Shabu Town Shibas
Shaggy Sheep	Shiba Society	Sidus NFT Heroes	SingularityHeroes	Sipherian Surge
Skvullpvks Hideout	Slacker Duck Pond	Sleeper Hits Collection Volume 1 NFT Cribs	Slimes World	Slumdoge Billionaires
Sneaky Vampire Syndicate	Soccer Doge Club	Space Dinos Club	Space Poggers	SpacePunksClub
SpaceShibas	Spookies NFT	SportsLeon Lion Club	Spunks	Standametti
Stoned Apez Saturn Club	Stoner Cats	Stranger EggZ	StripperVille NFTs	SupDucks
Super Yeti	Superfuzz The Bad Batch	Superfuzz The Good Guys	Sushiverse	SympathyForTheDevils
THE PLUTO ALLIANCE	THE SHRUNKENHEADZ	The Alien Boy	The BirdHouse	The CryptoDads
The CryptoSaints	The Doge Pound	The Fuckin' Trolls	The Goobers	The Graveyard Sale
The KILLAZ	The KittyButts	The League Of Sacred Devils	The Lost Glitches	The MonstroCities
The Moon Boyz	The NFTBirds	The Nanoz	The Nemesis Companions	The Ninja Hideout
The Project URS	The Sevens (Official)	The Shark Cove	The Soldiers Of The Metaverse	The Street Dawgs
The Unstable Horses Yard	The Vogt Collective	The Wanderers	The Wicked Craniums	The Wicked Stallions
The WolfGang Pups	The WonderQuest	The WynLambo	TheHeartProject	TheTigersGuild
Tie Dye Ninjas	Tokenmon	Tools of Rock	Top Dog Beach Club	TradeSquads
Trollz	Ugly Cuties Art Club (UCAC)	United Punks Union	Untamed Elephants	Unusual Whales
VeeFriends	Vegiemon	Vox Collectibles	Voxies	WE ARE THE OUTKAST
Waifusion	Wall Street Chads	Wanna Panda	Wannabes Music Club	Warriors of Aradena
We are Dorkis	WeMint Washington	Weird Whales	Wicked Ape Bone Club	Wicked Hound Bone Club
Wild Stag Treehouse	Winter Bears	Woodies Generative Characters	World of Women	Zunks
astroGems	bastard gan penguins	isotile Genesis Avatars	thedudes	uwucrew

## F.4 Rarity score distributions

As detailed in the main text, we use Akaike Information Criterion [2] and Maximum Likelihood Estimation to determine the distribution that best describes the rarity score distribution for each collection. We select the distribution among a subset of distributions implemented in the `scipy.stats` python package, requiring the distributions to be heterogeneous, continuous, and with at most 3 parameters (including location and scale). This results in choosing among the following distributions: *uniform*, *pareto*, *cauchy*, *lognormal*, *levy*, *exponential*.

## F.5 Additional figures




Figure F.1: **Collectible NFTs minted over time.** Distribution of the collectible NFTs considered in this analysis minted over time.




Figure F.2: **Distribution of the log-normal distribution characteristic parameter  $\mu$ .** Distribution of the log-normal distribution parameter  $\mu$  (blue line), and its average value across collections (red dashed line). The log-normal distribution  $\ln(X) \sim \mathcal{N}(\mu, \sigma^2)$  captures the distribution of rarity for 90% of collections.




Figure F.3: **Fitting of the rarity distribution.** Distribution of the rarity score of the NFTs within several collections included in the dataset (blue dots), along with the best distribution fit computed using Maximum Likelihood Estimation and Akaike Information Criterion [2] (orange line).

## F.5.1 Rarity rank results




Figure F.4: **Trait Rarity Rank.** a-c) Distribution of the trait rarity rank of the NFTs within three collections: CryptoPunks (a), Bored Ape Yacht Club (b), and World of Women (c). d) Violin plot of the Spearman Rank correlation computed between the rarity rank and the number of NFTs with that rank.




Figure F.5: **Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis with the rarity rank.** Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity  $(100-q)$ , where  $q$  is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. The NFTs are aggregated by quantile depending on their rarity rank, i.e the first quantile represents the rarest NFTs within the collection.




Figure F.6: **High rarity leads to higher returns, and a lower chance of a negative return - Analysis with the rarity rank.** a) Median return in USD by rarity quantile. Inset: median return against the quantity  $(100-q)$ , where  $q$  is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. The NFTs are aggregated by quantile depending on their rarity rank, i.e the first quantile represents the rarest NFTs within the collection.

## F.5.2 Currency robustness check




Figure F.7: **Rare NFTs have a higher financial value and circulate less on the marketplace - Price in ETH.** Median sale price in ETH (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity  $(100-q)$ , where  $q$  is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections.




Figure F.8: **High rarity leads to higher returns, and a lower chance of a negative return - Price in ETH.** a) Median return in ETH by rarity quantile. Inset: median return against the quantity  $(100-q)$ , where  $q$  is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in ETH by rarity quantile.

### F.5.3 Time robustness check

To make sure that the findings we highlight in this paper are time-independent, we ran the same analysis by using only the transactions happening during specific time periods, to see whether we observe the same mechanisms within the marketplace. Therefore, we performed the analysis on the two last quarters of 2021, i.e., first on Q3 (July - September 2021) and then on Q4 (October - December 2021).




Figure F.9: **Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis on Q3 2021.** Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity  $(100-q)$ , where  $q$  is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. This analysis only takes into consideration the sales happening during Q3 2021.




Figure F.10: **High rarity leads to higher returns, and a lower chance of a negative return - Analysis on Q3 2021.** a) Median return in USD by rarity quantile. Inset: median return against the quantity  $(10^2 - q)$ , where  $q$  is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. This analysis only takes into consideration the sales happening during Q3 2021.




Figure F.11: **Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis on Q4 2021.** Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity  $(10^2 - q)$ , where  $q$  is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. This analysis only takes into consideration the sales happening during Q4 2021.




Figure F.12: **High rarity leads to higher returns, and a lower chance of a negative return - Analysis on Q4 2021.** a) Median return in USD by rarity quantile. Inset: median return against the quantity  $(100-q)$ , where  $q$  is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. This analysis only takes into consideration the sales happening during Q4 2021.

## F.5.4 Tails robustness check




Figure F.13: **Rare NFTs have a higher financial value and circulate less on the marketplace - Analysis without the rarest and least rare NFTs.** Median sale price in USD (a-c) and average number of sales (e-g) by rarity quantile (with 20 quantiles considered) for three collections: CryptoPunks (a and e), Bored Ape Yacht Club (b and f), and World Women (c and g). d) Median sale price by rarity quantile (with 100 quantiles considered) considering all collections. Inset: median sale price against the quantity  $(10^2 - q)$ , where  $q$  is the rarity quantile, in log-log scale (black line) and the corresponding power law fit (green dashed line). h) Median number of sales by rarity quantile considering all collections. This analysis was performed after discarding the 10% rarest and least rare NFTs from each collection.




Figure F.14: **High rarity leads to higher returns, and a lower chance of a negative return - Analysis without the rarest and least rare NFTs.** a) Median return in USD by rarity quantile. Inset: median return against the quantity  $(100-q)$ , where  $q$  is the rarity quantile in log-log scale (black line) and the corresponding power law fit (green dashed line). b) Fraction of sales with negative return in USD by rarity quantile. This analysis was performed after discarding the 10% rarest and least rare NFTs from each collection.

# Bibliography

- [1] Matthieu Nadini, Christian Bongiorno, Alessandro Rizzo, and Maurizio Porfiri. Detecting network backbones against time variations in node properties. *Nonlinear Dynamics*, 99(1):855–878, 2020.
- [2] Eric-Jan Wagenmakers and Simon Farrell. Aic model selection using akaike weights. *Psychonomic bulletin & review*, 11(1):192–196, 2004.
- [3] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Technical report, Manubot, 2008.
- [4] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.
- [5] Amount of cryptocurrency held in decentralized finance, or defi, worldwide from august 2017 to january 12, 2022, 2022. Accessed Mar 9, 2022.
- [6] Defi market commentary | february 2022. <https://consensys.net/blog/cryptoeconomic-research/defi-market-commentary-february-2022/> Accessed September 19, 2022, 2022.
- [7] Dune analytics. <https://dune.xyz/rchen8/opensea> Accessed April 2, 2022, 2022.
- [8] Abeer ElBahrawy, Laura Alessandretti, Leonid Rusnac, Daniel Goldsmith, Alexander Teytelboym, and Andrea Baronchelli. Collective dynamics of dark web marketplaces. *Scientific reports*, 10(1):1–8, 2020.
- [9] Damilare Peter Oyinloye, Je Sen Teh, Norziana Jamil, and Moatsum Alawida. Blockchain consensus: An overview of alternative protocols. *Symmetry*, 13(8):1363, 2021.
- [10] Eric Ke Wang, RuiPei Sun, Chien-Ming Chen, Zuodong Liang, Saru Kumari, and Muhammad Khurram Khan. Proof of x-repute blockchain consensus protocol for iot systems. *Computers & Security*, 95:101871, 2020.
- [11] Abeer ElBahrawy, Laura Alessandretti, Anne Kandler, Romualdo Pastor-Satorras, and Andrea Baronchelli. Evolutionary dynamics of the cryptocurrency market. *Royal Society open science*, 4(11):170623, 2017.

- [12] Laura Alessandretti, Abeer ElBahrawy, Luca Maria Aiello, and Andrea Baronchelli. Anticipating cryptocurrency prices using machine learning. *Complexity*, 2018, 2018.
- [13] Ross C Phillips and Denise Gorse. Cryptocurrency price drivers: Wavelet coherence analysis revisited. *PloS one*, 13(4):e0195200, 2018.
- [14] Alberto Bracci, Matthieu Nadini, Maxwell Aliapoulios, Damon McCoy, Ian Gray, Alexander Teytelboym, Angela Gallo, and Andrea Baronchelli. Dark Web Marketplaces and COVID-19: Before the vaccine. *EPJ Data Science*, 10, 2021.
- [15] Monica J Barratt, Simon Lenton, Alexia Maddox, and Matthew Allen. “what if you live on top of a bakery and you like cakes?”—drug use and harm trajectories before, during and after the emergence of silk road. *International Journal of Drug Policy*, 35:50–57, 2016.
- [16] Roderic Broadhurst, Matthew Ball, and Chuxuan Jiang. Availability of COVID-19 related products on Tor darknet markets. *Australian Institute of Criminology*, April 2020.
- [17] Matthieu Nadini, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. Mapping the nft revolution: market trends, trade networks, and visual features. *Scientific reports*, 11(1):1–11, 2021.
- [18] Kishore Vasan, Milan Janosov, and Albert-László Barabási. Quantifying nft-driven networks in crypto art. *Scientific Reports*, 12, 02 2022.
- [19] Massimo Franceschet. Hits hits art. *Blockchain: Research and Applications*, 2(4):100038, 2021.
- [20] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, Rodney Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler Murray, Hsu-Han Ooi, Matthew Peters, Joanna Power, Sam Skjonsberg, Lucy Lu Wang, Chris Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. Construction of the literature graph in semantic scholar. In *NAACL*, 2018.
- [21] KR1442 Chowdhary. Natural language processing. *Fundamentals of artificial intelligence*, pages 603–649, 2020.
- [22] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019.
- [23] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

- [24] Alain Barrat, Marc Barthelemy, Romualdo Pastor-Satorras, and Alessandro Vespignani. The architecture of complex weighted networks. *Proceedings of the national academy of sciences*, 101(11):3747–3752, 2004.
- [25] Nicola Perra, Bruno Gonçalves, Romulo Pastor-Satorras, and Alessandro Vespignani. Activity driven modeling of time varying networks. *Scientific Reports*, 2:469, 2012.
- [26] Martin Gerlach, Tiago P. Peixoto, and Eduardo G. Altmann. A network approach to topic models. *Science Advances*, 4(7):eaq1360, 2018.
- [27] Ye Sun and Vito Latora. The evolution of knowledge within and across fields in modern physics. *Scientific Reports*, 10(1):12097, December 2020.
- [28] Matthieu Nadini, Alessandro Rizzo, and Maurizio Porfiri. Reconstructing irreducible links in temporal networks: which tool to choose depends on the network size. *Journal of Physics: Complexity*, 1(1):015001, 2020.
- [29] Abeer ElBahrawy, Laura Alessandretti, Leonid Rusnac, Daniel Goldsmith, Alexander Teytelboym, and Andrea Baronchelli. Collective dynamics of dark web marketplaces. *Scientific reports*, 10(1):1–8, 2020.
- [30] Flashpoint intelligence. <https://www.flashpoint-intel.com/> Accessed January 4, 2021, 2020.
- [31] Chainalysis, inc. <https://www.chainalysis.com/> Accessed October 27, 2021, 2020.
- [32] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press, 2016.
- [33] Stuart Haber and W Scott Stornetta. How to time-stamp a digital document. In *Conference on the Theory and Application of Cryptography*, pages 437–455. Springer, 1990.
- [34] J Lawrence Carter and Mark N Wegman. Universal classes of hash functions. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 106–112, 1977.
- [35] Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker. Search and replication in unstructured peer-to-peer networks. In *Proceedings of the 16th international conference on Supercomputing*, pages 84–95, 2002.
- [36] Carlo Campajola, Raffaele Cristodaro, Francesco Maria De Collibus, Tao Yan, Nicolo’ Vallarano, and Claudio J Tessone. The evolution of centralisation on cryptocurrency platforms. *arXiv preprint arXiv:2206.05081*, 2022.
- [37] Jordi Herrera-Joancomartí. Research and challenges on bitcoin anonymity. In *Data privacy management, autonomous spontaneous security, and security assurance*, pages

- 3–16. Springer, 2014.
- [38] Arthur Gervais, Ghassan O Karame, Vedran Capkun, and Srdjan Capkun. Is bitcoin a decentralized currency? *IEEE security & privacy*, 12(3):54–60, 2014.
- [39] Ayana Aspembitova, Ling Feng, Valentin Melnikov, and Lock Yue Chew. Fitness preferential attachment as a driving mechanism in bitcoin transaction network. *PloS one*, 14(8):e0219346, 2019.
- [40] Elli Androulaki, Ghassan O Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. Evaluating user privacy in Bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 34–51. Springer, 2013.
- [41] Seunghyeon Lee, Changhoon Yoon, Heedo Kang, Yeonkeun Kim, Yongdae Kim, Dongsu Han, Sooel Son, and Seungwon Shin. Cybercriminal minds: An investigative study of cryptocurrency abuses in the dark web. In *Network and Distributed System Security Symposium*, pages 1–15. Internet Society, 2019.
- [42] Sean Foley, Jonathan R Karlsen, and Tālis J Putniņš. Sex, drugs, and Bitcoin: How much illegal activity is financed through cryptocurrencies? *The Review of Financial Studies*, 32(5):1798–1853, 2019.
- [43] John Markoff. *What the dormouse said: How the sixties counterculture shaped the personal computer industry*. Penguin, New York, USA, 2005.
- [44] Tim Berners-Lee, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret. The world-wide web. *Communications of the ACM*, 37(8):76–82, 1994.
- [45] Mike Power. Drugs 2.0: The web revolution that’s changing how the world gets high. <https://www.theguardian.com/books/2013/may/12/drugs-web-revolution-power-review> Accessed January 4, 2021, 2013. The Guardian.
- [46] Kim Zetter. 8 suspects arrested in online drug market sting. <https://www.wired.com/2012/04/online-drug-market-takedown/> Accessed January 4, 2021, 2012. Wired.
- [47] Monica J Barratt. Discussing illicit drugs in public internet forums: visibility, stigma, and pseudonymity. In *Proceedings of the 5th International Conference on Communities and Technologies*, pages 159–168, 2011.
- [48] James Martin. *Drugs on the dark net: How cryptomarkets are transforming the global trade in illicit drugs*. Springer, London, UK, 2014.
- [49] Nicolas Christin. 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, 2013.

- [50] Monica J Barratt, Jason A Ferris, and Adam R Winstock. Use of Silk Road, the online drug marketplace, in the United Kingdom, Australia and the United States. *Addiction*, 109(5):774–783, 2014.
- [51] James Martin. Lost on the Silk Road: Online drug distribution and the cryptomarket. *Criminology & Criminal Justice*, 14(3):351–367, 2014.
- [52] Judith Aldridge and David Décary-Héту. Not an Ebay for drugs: The cryptomarket Silk Road as a paradigm shifting criminal innovation. *Available at SSRN 2436643*, 2014.
- [53] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, Naval Research Lab Washington DC, 2004.
- [54] Francesco Cappa and Michele Pinelli. Collecting money through blockchain technologies: first insights on the determinants of the return on initial coin offerings. *Information Technology for Development*, pages 1–18, 2020.
- [55] Marco Masoni, Maria Renza Guelfi, and Gian Franco Gensini. Darknet and Bitcoin, the obscure and anonymous side of the internet in healthcare. *Technology and Health Care*, 24(6):969–972, 2016.
- [56] Damien Rhumorbarbe, Ludovic Staehli, Julian Broséus, Quentin Rossy, and Pierre Esseiva. Buying drugs on a darknet market: A better deal? studying the online illicit drug market through the analysis of digital, physical and chemical data. *Forensic science international*, 267:173–182, 2016.
- [57] Gerard E. Lynch. United States of America : Vs. Ross William Ulbricht. <https://caselaw.findlaw.com/us-2nd-circuit/1862572.html> Accessed January 4, 2021, 2017. United States Court of Appeals.
- [58] Gwern. Darknet market mortality risks. <https://www.gwern.net/DNM-survival> Accessed October 27, 2021, 2019.
- [59] Marie Claire Van Hout and Tim Bingham. Silk Road, the virtual drug marketplace: A single case study of user experiences. *International Journal of Drug Policy*, 24(5):385–391, 2013.
- [60] Marie Claire Van Hout and Tim Bingham. Responsible vendors, intelligent consumers: Silk Road, the online revolution in drug trading. *International Journal of Drug Policy*, 25(2):183–189, 2014.
- [61] Bassam Zantout, Ramzi Haraty, et al. I2p data communication system. In *Proceedings of ICN*, pages 401–409. Citeseer, 2011.
- [62] Frank Wehinger. The dark net: Self-regulation dynamics of illegal online markets for identities and related services. In *2011 European Intelligence and Security Informatics Conference*, pages 209–213. IEEE, 2011.

- [63] Darren Hayes, Francesco Cappa, and Nhien An Le-Khac. An effective approach to mobile device management: Security and privacy issues associated with mobile applications. *Digital Business*, 1(1):100001, 2020.
- [64] Malte Möser, Kyle Soska, Ethan Heilman, Kevin Lee, Henry Heffan, Shashvat Srivastava, Kyle Hogan, Jason Hennessey, Andrew Miller, Arvind Narayanan, et al. An empirical analysis of traceability in the monero blockchain. *Proceedings on Privacy Enhancing Technologies*, 2018(3):143–163, 2018.
- [65] Kyle Soska and Nicolas Christin. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *24th {USENIX} security symposium ({USENIX} security 15)*, pages 33–48, 2015.
- [66] Nicolas Christin. An eu-focused analysis of drug supply on the alphabay marketplace. *EMCDDA commissioned paper.*, 2017.
- [67] Ionut Ilascu. Dark web’s wall street market and valhalla seized, six arrested. <https://www.bleepingcomputer.com/news/security/dark-web-s-wall-street-market-and-valhalla-seized-six-arrested/> Accessed January 4, 2021, 2019. Bleeping Computer.
- [68] Darknet takedown: Authorities shutter online criminal market AlphaBay. <https://www.fbi.gov/news/stories/alphabay-takedown> Accessed October 27, 2021, 2017. FBI.
- [69] Grey market unsurprisingly exit scammed. <https://darknetlive.com/post/grey-market-unsurprisingly-exit-scammed/> Accessed January 4, 2021, 2019. DarknetLive.
- [70] Bitbazaar is exit scamming. <https://darknetlive.com/post/bitbazaar-is-exit-scamming/> Accessed January 4, 2021, 2019. DarknetLive.
- [71] Dread forum. <https://onion.live/site/dread-forum> Accessed via Tor browser October 27, 2021, 2020.
- [72] Kyle Porter. Analyzing the darknetmarkets subreddit for evolutions of tools and trends using lda topic modeling. *Digital Investigation*, 26:S87 – S97, 2018.
- [73] Onion.live. <https://onion.live/page/about> Accessed via Tor browser August 2, 2020, 2020.
- [74] Darknetlive. <https://darknetlive.com/> Accessed October 27, 2021, 2020.
- [75] Dark.fail. <https://dark.fail/> Accessed October 27, 2021, 2020.
- [76] Dnstats. <https://dnstats.net/> Accessed January 4, 2021, 2020.
- [77] Wesley Lacson and Beata Jones. The 21st century darknet market: Lessons from the fall of silk road. *International Journal of Cyber Criminology*, 10(1), 2016.

- [78] Diana S Dolliver and Katherine L Love. Criminogenic asymmetries in cyberspace: A comparative analysis of two tor marketplaces. *Journal of Globalization Studies*, 6(2):75–96, 2015.
- [79] Diana S Dolliver. Evaluating drug trafficking on the Tor network: Silk Road 2, the sequel. *International Journal of Drug Policy*, 26(11):1113–1123, 2015.
- [80] Julian Broséus, Damien Rhumorbarbe, Caroline Mireault, Vincent Ouellette, Frank Crispino, and David Décary-Hétu. Studying illicit drug trafficking on darknet markets: Structure and organisation from a Canadian perspective. *Forensic Science International*, 264:7–14, 2016.
- [81] Rolf van Wegberg, Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Hernandez Ganán, Bram Klievink, Nicolas Christin, and Michel van Eeten. Plug and prey? measuring the commoditization of cybercrime via online anonymous markets. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1009–1026, Baltimore, MD, August 2018. USENIX Association.
- [82] Muhammad Izhar Mehar, Charles Louis Shier, Alana Giambattista, Elgar Gong, Gabrielle Fletcher, Ryan Sanayhie, Henry M Kim, and Marek Laskowski. Understanding a revolutionary and flawed grand experiment in blockchain: the dao attack. *Journal of Cases on Information Technology (JCIT)*, 21(1):19–32, 2019.
- [83] What is ethereum? <https://ethereum.org/en/what-is-ethereum/> Accessed September 19, 2022, 2022.
- [84] Beeple sold an nft for \$69 million. <https://www.theverge.com/2021/3/11/22325054/beeple-christies-nft-sale-cost-everydays-69-million> Accessed September 19, 2022, 2022.
- [85] Arnav Kapoor, Dipanwita Guhathakurta, Mehul Mathur, Rupanshu Yadav, Manish Gupta, and Ponnurangam Kumaraguru. Tweetboost: Influence of social media on nft valuation. In *Companion Proceedings of the Web Conference 2022, WWW '22*, page 621–629, New York, NY, USA, 2022. Association for Computing Machinery.
- [86] Sebeom Oh, Samuel Rosen, and Anthony Lee Zhang. Investor experience matters: Evidence from generative art collections on the blockchain. *Available at SSRN 4042901*, 2022.
- [87] Hong Bao and David Roubaud. Recent development in fintech: Non-fungible token, 2021.
- [88] Victor von Wachter, Johannes Rude Jensen, Ferdinand Regner, and Omri Ross. Nft wash trading: Quantifying suspicious behaviour in nft markets. In *Financial Cryptography and Data Security. FC 2022 International Workshops*, 2021.
- [89] Konstantinos Pelechrinis, Xin Liu, Prashant Krishnamurthy, and Amy Babay. Spot-

- ting anomalous trades in nft markets: The case of nba topshot. *arXiv preprint arXiv:2202.04013*, 2022.
- [90] Lennart Ante. The non-fungible token (nft) market and its relationship with bitcoin and ethereum. *Available at SSRN 3861106*, 2021.
- [91] Nicola Borri, Yukun Liu, and Aleh Tsyvinski. The economics of non-fungible tokens. *Available at SSRN*, 2022.
- [92] David Y Aharon and Ender Demir. Nfts and asset class spillovers: Lessons from the period around the covid-19 pandemic. *Finance Research Letters*, page 102515, 2021.
- [93] Mark Zuckerberg is betting Facebook’s future on the metaverse. <https://www.theverge.com/22588022/mark-zuckerberg-facebook-ceo-metaverse-interview> Accessed March 19, 2022, 2021.
- [94] Bored Ape Yacht Club creator raises \$450 million to build an NFT metaverse. <https://www.theverge.com/2022/3/22/22991272/yuga-labs-seed-funding-a16z-bored-ape-yacht-club-bayc-metaverse-other-side> Accessed March 23, 2022, 2022.
- [95] Qiinglin Ynag, Yetong Zhao, Huawei Huang, and Zibin Zheng. Fusing blockchain and ai with metaverse: A survey. *arXiv preprint arXiv:2201.03201*, 2022.
- [96] Lik-Hang Lee, Zijun Lin, Rui Hu, Zhengya Gong, Abhishek Kumar, Tangyao Li, Sijia Li, and Pan Hui. When creators meet the metaverse: A survey on computational arts. *arXiv preprint arXiv:2111.13486*, 2021.
- [97] J. Broséus, D. Rhumorbarbe, C. Mireault, V. Ouellette, F. Crispino, and D. Décary-Hétu. Studying illicit drug trafficking on Darknet markets: Structure and organisation from a Canadian perspective. *Forensic Science International*, 264:7–14, July 2016.
- [98] Jack Cunliffe, James Martin, David Décary-Hétu, and Judith Aldridge. An island apart? risks and prices in the Australian cryptomarket drug trade. *International Journal of Drug Policy*, 50:64–73, December 2017.
- [99] Pierluigi Paganini. Italian police shut down darkweb berlusconi market and arrested admins. <https://securityaffairs.co/wordpress/93603/cyber-crime/berlusconi-market-darkweb.html> Accessed January 4, 2021, 2019. Security Affairs.
- [100] Europol. Operation Onymous. <https://www.europol.europa.eu/activities-services/europol-in-action/operations/operation-onymous> Accessed October 27, 2021, 2014.
- [101] Julia Buxton and Tim Bingham. The rise and challenge of dark net drug markets. *Policy brief*, 7:1–24, 2015.

- [102] Kyle Soska and Nicolas Christin. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *Proceedings of the 24th USENIX Conference on Security Symposium, SEC'15*, pages 33–48, USA, 2015. USENIX Association.
- [103] Andres Baravalle, Mauro Sanchez Lopez, and Sin Wee Lee. Mining the Dark Web: Drugs and Fake Ids. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, pages 350–356, Barcelona, Spain, December 2016. IEEE.
- [104] Nicolas Christin. Traveling the silk road: A measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22nd International Conference on World Wide Web - WWW '13*, pages 213–224, Rio de Janeiro, Brazil, 2013. ACM Press.
- [105] Judith Aldridge and David DDcary-HHtu. Not an 'Ebay for Drugs': The Cryptomarket 'Silk Road' as a Paradigm Shifting Criminal Innovation. *SSRN Electronic Journal*, 2014.
- [106] Rolf Van Wegberg, Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Gañán, Bram Klievink, Nicolas Christin, and Michel Van Eeten. Plug and Prey? measuring the Commoditization of Cybercrime via Online Anonymous Markets. In *Proceedings of the 27th USENIX Conference on Security Symposium, SEC'18*, pages 1009–1026, USA, 2018. USENIX Association.
- [107] Lukas Norbutas. Offline constraints in online drug marketplaces: An exploratory analysis of a cryptomarket trade network. *International Journal of Drug Policy*, 56:92–100, June 2018.
- [108] Martin Dittus, Joss Wright, and Mark Graham. Platform Criminalism: The 'Last-Mile' Geography of the Darknet Market Supply Chain. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*, pages 277–286, Lyon, France, 2018. ACM Press.
- [109] Bitcoincore. <https://bitcoin.org/en/bitcoin-core/> Accessed October 27, 2021, 2020.
- [110] Blockchain.com. [www.blockchain.com](http://www.blockchain.com) Accessed October 27, 2021, 2020. Daily Mail.
- [111] Dorit Ron and Adi Shamir. Quantitative analysis of the full Bitcoin transaction graph. In *International Conference on Financial Cryptography and Data Security*, pages 6–24. Springer, 2013.
- [112] Paolo Tasca, Adam Hayes, and Shaowen Liu. The evolution of the Bitcoin economy. *The Journal of Risk Finance*, 2018.
- [113] Martin Harrigan and Christoph Fretter. The unreasonable effectiveness of address clustering. In *2016 Intl IEEE Conferences on Ubiquitous Intelligence & Computing, Advanced and Trusted Computing, Scalable Computing and Communications*,

- Cloud and Big Data Computing, Internet of People, and Smart World Congress (UIC/ATC/ScalCom/CBDCCom/IoP/SmartWorld)*, pages 368–373. IEEE, 2016.
- [114] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. A fistful of Bitcoins: Characterizing payments among men with no names. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 127–140, 2013.
- [115] Wikipedia now accepts bitcoin donations. <https://www.coindesk.com/wikipedia-now-accepts-bitcoin-donations> Accessed October 27, 2021, 2014. Coindesk.
- [116] YoonJae Chung. Cracking the code: How the us government tracks Bitcoin transactions. *Analysis Of Applied Mathematics*, page 152, 2019.
- [117] Chainalysis. Chainalysis in action: How law enforcement used blockchain analysis to follow funds and identify the twitter hackers. <https://blog.chainalysis.com/reports/chainalysis-doj-twitter-hack-2020> Accessed October 27, 2021, 2020.
- [118] Lily Hay Newman. How a Bitcoin Trail Led to a Massive Dark Web Child-Porn Site Takedown, 2019.
- [119] N. de la Rouviere. A subgraph to index & explore CryptoKitties auctions. <https://thegraph.com/explorer/subgraph/nieldlr/cryptokitties-sales>. (Accessed 4 May 2021), 2021.
- [120] D. Rosenbaum. Gods-Unchained marketplace. <https://github.com/djrosenbaum/unchained-transactions>. (Accessed 4 May 2021), 2021.
- [121] D. Team. Decentraland marketplace. <https://thegraph.com/explorer/subgraph/decentraland/marketplace>. (Accessed 4 May 2021), 2021.
- [122] Opensea api overview. <https://docs.opensea.io/reference/api-overview> Accessed June 28, 2022.
- [123] N. Team. The best place to analyze, track, and discover NFTs. <https://nonfungible.com/>. (Accessed 4 May 2021), 2021.
- [124] Balázs Bodó, Alexandra Giannopoulou, João Quintais, and Péter Mezei. The rise of nfts: These aren’t the droids you’re looking for. 2022.
- [125] Derrick Bryson Taylor. A timeline of the coronavirus pandemic. <https://www.nytimes.com/article/coronavirus-timeline.html> Accessed January 4, 2021, 2020. The New York Times.
- [126] Moritz UG Kraemer, Chia-Hung Yang, Bernardo Gutierrez, Chieh-Hsi Wu, Brennan Klein, David M Pigott, Louis Du Plessis, Nuno R Faria, Ruoran Li, William P Hanage,

- et al. The effect of human mobility and control measures on the COVID-19 epidemic in china. *Science*, 368(6490):493–497, 2020.
- [127] Nuno Fernandes. Economic effects of coronavirus outbreak (COVID-19) on the world economy. *Available at SSRN 3557504*, 2020.
- [128] Rabah Arezki and Ha Nguyen. 4 novel coronavirus hurts the middle east and north africa through many channels. *Economics in the Time of COVID-19*, page 53, 2020.
- [129] Markets energy. <https://www.bloomberg.com/energy> Accessed January 4, 2021, 2020. Bloomberg.
- [130] Abby Narishkin and Steve Cameron. Why toilet-paper demand spiked 845%, and how companies kept up with it. <https://www.businessinsider.com/why-toilet-paper-demand-spiked-845-how-companies-kept-up-2020-5?r=US&IR=T> Accessed January 4, 2021, 2020.
- [131] Anna Medaris Miller. The Wuhan coronavirus has led to a face mask shortage, with sellersnow offering masks at up to \$7 apiece. <https://www.businessinsider.com/wuhan-coronavirus-mask-shortage-causing-price-increase-2020-1?IR=T> Accessed January 4, 2021, 2020. Business Insider.
- [132] Kazuki Sakuma. Woman reported to prosecutors for alcohol disinfectant at double price in japan. <https://mainichi.jp/english/articles/20200701/p2a/00m/0na/012000c> Accessed January 4, 2021, 2020. The Mainichi.
- [133] Rik Chakraborti and Gavin Roberts. Anti-price gouging laws, shortages, and COVID-19: Big data insights from consumer searches.
- [134] GIATOC. CovidCrimeWatch: crime and contagion: the impact of a pandemic on organized crime. <https://globalinitiative.net/crime-contagion-impact-covid-crime/> Accessed January 4, 2021, 2020.
- [135] Europol and EMCDDA. EU drug markets impact of COVID-19. <https://www.europol.europa.eu/newsroom/news/eu-drug-markets-impact-of-covid-19> Accessed January 4, 2021, 2020.
- [136] Covid-19 vaccine: first person receives Pfizer jab in UK. <https://www.bbc.co.uk/news/uk-55227325> Accessed December 8, 2020, 2020. BBC.
- [137] WHO. Coronavirus disease (covid-19) advice for the public: when and how to use masks. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/advice-for-public/when-and-how-to-use-masks> Accessed January 4, 2021, 2020. World Health Organization.
- [138] Coronavirus disease 2019 (covid-19): Emerging treatments. <https://bestpractice.bmj.com/topics/en-gb/3000168/emergingtxs> Accessed January 4, 2021, 2020.

- [139] Emily Chen, Kristina Lerman, and Emilio Ferrara. Tracking social media discourse about the COVID-19 pandemic: Development of a public coronavirus Twitter data set. *JMIR Public Health and Surveillance*, 6(2):e19273, 2020.
- [140] Wikipedia. Wikimedia REST API. [https://wikimedia.org/api/rest\\_v1/](https://wikimedia.org/api/rest_v1/) Accessed January 4, 2021, 2020.
- [141] Matthew Ball, Roderic Broadhurst, Alexander Niven, and Harshit Trivedi. Data capture and analysis of darknet markets. *Available at SSRN 3344936*, 2019.
- [142] Andres Baravalle, Mauro Sanchez Lopez, and Sin Wee Lee. Mining the dark web: Drugs and fake ids. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, pages 350–356. IEEE, 2016.
- [143] Alessandro Celestini, Gianluigi Me, and Mara Mignone. Tor marketplaces exploratory data analysis: The drugs case. In *International Conference on Global Security, Safety, and Sustainability*, pages 218–229. Springer, 2017.
- [144] Darren R Hayes, Francesco Cappa, and James Cardon. A framework for more effective dark web marketplace investigations. *Information*, 9(8):186, 2018.
- [145] D Décarry-Héту and J Aldridge. Datacrypto: The dark net crawler and scraper. *Software Program*, 2013.
- [146] X-byte enterprise crawling. <https://www.xbyte.io/dark-deep-web-data-scraping.php> Accessed January 4, 2021.
- [147] Dark net markets comparison chart. <https://www.darknetstats.com/dark-net-markets-comparison-chart/> Accessed January 4, 2021, 2020.
- [148] James Gallagher. Dexamethasone, remdesivir, regeneron: Trump’s covid treatment explained. <https://www.bbc.co.uk/news/health-54418464> Accessed December 8, 2020, 2020. BBC.
- [149] Jason Horowitz and Emma Bubola. On day 1 of lockdown, italian officials urge citizens to abide by rules. <https://www.nytimes.com/2020/03/08/world/europe/italy-coronavirus-quarantine.html> Accessed January 4, 2021, 2020. The New York Times.
- [150] Mark Landler and Stephen Castle. Britain locks down to stem the coronavirus. more or less. <https://www.nytimes.com/2020/03/24/world/europe/britain-coronavirus-lockdown.html> Accessed January 4, 2021, 2020. The New York Times.
- [151] Marc Santora. The world reopens, despite skyrocketing coronavirus cases. <https://www.nytimes.com/2020/06/09/world/coronavirus-reopenings.html> Accessed January 4, 2021, 2020. The New York Times.

- [152] Sam Jones, Kim Willsher, and Natalie Grover. Spain is first country in western Europe to record half a million COVID cases. <https://www.theguardian.com/world/2020/sep/07/spain-set-to-record-half-a-million-covid-cases> Accessed January 24, 2021, 2020. The Guardian.
- [153] COVID-19: Pm announces four-week england lockdown. <https://www.bbc.co.uk/news/uk-54763956> Accessed November 28, 2020, 2020. BBC.
- [154] Remarks by president trump, vice president pence, and members of the coronavirus task force in press briefing. <https://www.whitehouse.gov/briefings-statements/remarks-president-trump-vice-president-pence-members-coronavirus-task-force-press-> Accessed January 4, 2021, 2020. White House.
- [155] Remarks by president trump at signing of h.r.748, the cares act. <https://www.whitehouse.gov/briefings-statements/remarks-president-trump-signing-h-r-748-cares-act/> Accessed November 27, 2020, 2020. The New York Times.
- [156] Zachary Warmbrodt. Senate agrees to extend small business rescue in surprise move. <https://www.politico.com/news/2020/06/30/senate-small-business-346412> Accessed November 27, 2020, 2020. Politico.com.
- [157] Li Zhou. The senate failed to pass more stimulus for a struggling economy. here's why. <https://www.vox.com/2020/9/10/21429678/senate-stimulus-vote> Accessed November 30, 2020, 2020. Vox.
- [158] Nicolò Gozzi, Michele Tizzani, Michele Starnini, Fabio Ciulla, Daniela Paolotti, André Panisson, and Nicola Perra. Collective response to media coverage of the COVID-19 pandemic on Reddit and Wikipedia: mixed-methods analysis. *Journal of medical Internet research*, 22(10):e21597, 2020.
- [159] Riccardo Gallotti, Francesco Valle, Nicola Castaldo, Pierluigi Sacco, and Manlio De Domenico. Assessing the risks of “infodemics” in response to COVID-19 epidemics. *ArXiv:2004.03997*, 2020.
- [160] Matteo Cinelli, Walter Quattrociochi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoli, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. The COVID-19 social media infodemic. *arXiv preprint arXiv:2003.05004*, 2020.
- [161] Covid-19: Europe sees hope in decline in new infections. <https://www.aa.com.tr/en/europe/covid-19-europe-sees-hope-in-decline-in-new-infections/1817159> Accessed January 4, 2021, 2020. AA.
- [162] Ignoring expert opinion, trump again promotes use of hydroxychloroquine. <https://www.nytimes.com/2020/04/05/us/politics/trump-hydroxychloroquine-coronavirus.html> Accessed January 4, 2021, 2020. The New York Times.

- [163] Trump is taking hydroxychloroquine, white house confirms. <https://www.theguardian.com/us-news/2020/may/19/trump-hydroxychloroquine-covid-19-white-house> Accessed January 4, 2021, 2020. Guardian.
- [164] Joan E Greve and Martin Pengelly. Twitter limits donald trump jr’s account for posting COVID-19 misinformation. <https://www.theguardian.com/us-news/2020/jul/28/donald-trump-jr-twitter-restricted-hydroxychloroquine> Accessed November 28, 2020, 2020. The Guardian.
- [165] Coronavirus: confirmed global cases pass one million. <https://www.bbc.co.uk/news/world-52144390> Accessed January 4, 2021, 2020. BBC.
- [166] Chinmayi Arun. On whatsapp, rumours, and lynchings. *Economic & Political Weekly*, 54(6):30–35, 2019.
- [167] Onion.live: Darkbay’s trust report. <https://onion.live/reports/2005451?page=1> Accessed via Tor browser August 2, 2020, 2020.
- [168] Dnstats: Darkbay. <https://dnstats.net/site/darkbay> Accessed January 4, 2021, 2020.
- [169] Comparison uptime dark markets. <https://dnstats.net/markets/1> Accessed January 4, 2021, 2020.
- [170] Sam Miller, Abeer El-Bahrawy, Martin Dittus, Mark Graham, and Joss Wright. Predicting drug demand with wikipedia views: Evidence from darknet markets. In *Proceedings of The Web Conference 2020*, pages 2669–2675, 2020.
- [171] Angus Bancroft. *The darknet and smarter crime: methods for investigating criminal entrepreneurs and the illicit drug economy*. Springer Nature, Cham, Switzerland, 2019.
- [172] The COVID-19 candidate vaccine landscape. <https://www.who.int/publications/m/item/draft-landscape-of-covid-19-candidate-vaccines> Accessed January 21, 2021, 2021. WHO.
- [173] Talha Khan Burki. The russian vaccine for COVID-19. *The Lancet Respiratory Medicine*, 8(11):e85–e86, 2020.
- [174] Heidi Ledford, David Cyranoski, and Richard Van Noorden. The UK has approved a COVID vaccine-here’s what scientists now want to know. *Nature*, 588(7837):205–206, 2020.
- [175] Janice Hopkins Tanne. COVID-19: FDA panel votes to authorise Pfizer Biontech vaccine. *BMJ*, 371, 2020.
- [176] Roberts Michelle. COVID: Pfizer-Biontech vaccine approved for eu states. <https://www.bbc.co.uk/news/world-europe-55401136> Accessed January 22, 2021, 2020.

- BBC.
- [177] Ewen Callaway. The unequal scramble for coronavirus vaccines — by the numbers. *Nature*, 584(7822):506–507, August 2020.
- [178] Interpol warns of organized crime threat to COVID-19 vaccines. <https://www.interpol.int/en/News-and-Events/News/2020/INTERPOL-warns-of-organized-crime-threat-to-COVID-19-vaccines> Accessed January 22, 2021, 2020. Interpol.
- [179] Europol predictions correct for fake COVID-19 vaccines. <https://www.europol.europa.eu/newsroom/news/europol-predictions-correct-for-fake-covid-19-vaccines> Accessed January 22, 2021, 2020. EUROPOL.
- [180] Jamie Redman. After empire’s exit scam, darknet market patrons scramble to find alternatives. <https://news.bitcoin.com/after-empires-exit-scam-darknet-market-patrons-scramble-to-find-alternatives/> Accessed February 6, 2021, 2020. Bitcoin.com.
- [181] Kevin Helms. Darknet giant darkmarket shut down, alleged operator arrested. <https://news.bitcoin.com/darknet-darkmarket-shut-down-operator-arrested/#:~:text=Darknet%20Giant%20Darkmarket%20Shut%20Down,%20Alleged%20Operator%20Arrested,was%20arrested%20in%20Germany.%20Authorities%20Shut%20Down%20Darkmarket> Accessed October 27, 2021, 2021. Bitcoin.com.
- [182] Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *Journal of Open Source Software*, 2(11):205, 2017.
- [183] Lisa A Jackson, Evan J Anderson, Nadine G Roupheal, Paul C Roberts, Mamodikoe Makhene, Rhea N Coler, Michele P McCullough, James D Chappell, Mark R Denison, Laura J Stevens, et al. An mRNA vaccine against SARS-CoV-2—preliminary report. *New England Journal of Medicine*, 2020.
- [184] Oxford University/AstraZeneca COVID-19 vaccine approved. <https://www.gov.uk/government/news/oxford-universityastrazeneca-covid-19-vaccine-approved> Accessed January 24, 2021, 2020. UK press release.
- [185] Aliss Higham. Will i need proof of a coronavirus vaccine if i want to go on holiday this summer? <https://www.express.co.uk/travel/articles/1387164/Will-I-need-proof-of-coronavirus-vaccine-to-go-on-holiday-summer-2021-EVG> Accessed January 26, 2021, 2021. Express.
- [186] Rebecca CH Brown, Julian Savulescu, Bridget Williams, and Dominic Wilkinson. Passport to freedom? immunity passports for covid-19. *Journal of Medical Ethics*, 46(10):652–659, 2020.

- [187] Ema advises against use of ivermectin for the prevention or treatment of covid-19 outside randomised clinical trials. <https://www.ema.europa.eu/en/news/ema-advises-against-use-ivermectin-prevention-treatment-covid-19-outside-randomised-clinical-trials> Accessed August 10, 2021, 2021. EMA.
- [188] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [189] Andréanne Bergeron, David Décary-Héту, and Luca Giommoni. Preliminary findings of the impact of covid-19 on drugs crypto markets. *International Journal of Drug Policy*, 83:102870, 2020.
- [190] 2020 united states postal service crisis. [https://en.wikipedia.org/wiki/2020\\_United\\_States\\_Postal\\_Service\\_crisis](https://en.wikipedia.org/wiki/2020_United_States_Postal_Service_crisis) Accessed August 9, 2021, 2021. Wikipedia.
- [191] Grace Kay. COVID-19 vaccines are selling on the dark web for as much as \$ 1,000 in Bitcoin. <https://uk.news.yahoo.com/online-scammers-trying-sell-covid-160008180.html> Accessed January 24, 2021, 2021. Business Insider.
- [192] Pieter Arntz. Buying COVID-19 vaccines from the dark web? No thanks! <https://blog.malwarebytes.com/scams/2020/12/buying-covid-19-vaccines-from-the-darkweb/> Accessed January 24, 2021, 2020. Malwarebytes.
- [193] Isobel Asher Hamilton. Dark web vendors are selling shady coronavirus ‘vaccines’ for \$ 300, and there’s been an uptick in listings since the FDA authorized Pfizer’s shot. <https://www.businessinsider.com/covid-vaccines-for-sale-dark-web-coronavirus-shot-2020-12?r=US&IR=T> Accessed January 24, 2021, 2021. Business Insider.
- [194] Angela Giuffrida. Italy imposes “green pass” restrictions on unvaccinated people. <https://www.theguardian.com/world/2021/jul/22/italy-covid-19-green-pass-vaccinations-restrictions> Accessed July 7, 2021, 2021. Guardian.
- [195] Solange Mougin. French parliament approves law requiring covid pass for restaurants, travel starting in august. <https://www.france24.com/en/france/20210725-french-parliament-approves-law-requiring-restaurant-covid-pass-vaccine-rules> 2021. France24.
- [196] James Martin, Jack Cunliffe, David Décary-Héту, and Judith Aldridge. Effect of restricting the legal supply of prescription opioids on buying through online illicit marketplaces: interrupted time series analysis. *bmj*, 361, 2018.

- [197] Scott E Hadland and Leo Beletsky. Tighter prescribing regulations drive illicit opioid sales, 2018.
- [198] Matthieu Nadini, Alberto Bracci, Abeer ElBahrawy, Philip Gradwell, Alexander Teytelboym, and Andrea Baronchelli. Emergence and structure of decentralised trade networks around dark web marketplaces. *Scientific Reports*, 12(1):1–9, 2022.
- [199] Abeer ElBahrawy, Laura Alessandretti, Leonid Rusnac, Daniel Goldsmith, Alexander Teytelboym, and Andrea Baronchelli. Collective dynamics of dark web marketplaces. *Scientific Reports*, 10(1):18827, December 2020.
- [200] D. Décarry-Héту and L. Giommoni. Do police crackdowns disrupt drug cryptomarkets? a longitudinal analysis of the effects of Operation Onymous. *Crime, Law and Social Change*, 67(1):55–75, February 2017.
- [201] Alberto Bracci, Matthieu Nadini, Maxwell Aliapoulos, Damon McCoy, Ian Gray, Alexander Teytelboym, Angela Gallo, and Andrea Baronchelli. Dark Web Marketplaces and COVID-19: Before the vaccine. *EPJ data science*, 10(1):6, 2021.
- [202] Stephanie Chevalier. Annual net revenue of Amazon from 2004 to 2020, 2021. (accessed 2021-09-29).
- [203] Susan Standing, Craig Standing, and Peter E. D Love. A review of research on e-marketplaces 1997–2008. *Decision Support Systems*, 49(1):41–51, April 2010.
- [204] J. Yannis Bakos. Reducing Buyer Search Costs: Implications for Electronic Marketplaces. *Management Science*, 43(12):1676–1692, December 1997.
- [205] Vivek Choudhury, Kathleen S. Hartzel, and Benn R. Konsynski. Uses and Consequences of Electronic Markets: An Empirical Investigation in the Aircraft Parts Industry. *MIS Quarterly*, 22(4):471–507, 1998.
- [206] Nuray Terzi. The impact of e-commerce on international trade and employment. In *Encyclopedia of E-Commerce Development, Implementation, and Management*, pages 2271–2287. IGI Global, 2016.
- [207] James Martin. *Drugs on the Dark Net: How Cryptomarkets Are Transforming the Global Trade in Illicit Drugs*. Springer, 2014.
- [208] Matthew Elliott and Andrea Galeotti. The role of networks in antitrust investigations. *Oxford Review of Economic Policy*, 35(4):614–637, December 2019.
- [209] Itay P. Fainmesser and Andrea Galeotti. Pricing Network Effects. *The Review of Economic Studies*, 83(1):165–198, January 2016.
- [210] Mark Armstrong. Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691, September 2006.

- [211] Attila Ambrus and Rossella Argenziano. Asymmetric Networks in Two-Sided Markets. *American Economic Journal: Microeconomics*, 1(1):17–52, January 2009.
- [212] Geoffrey G. Parker and Marshall W. Van Alstyne. Two-Sided Network Effects: A Theory of Information Product Design. *Management Science*, 51(10):1494–1504, October 2005.
- [213] Chrysanthos Dellarocas. The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science*, 49(10):1407–1424, October 2003.
- [214] Thomas Eisenmann, Geoffrey Parker, and Marshall W Van Alstyne. Strategies for two-sided markets. *Harvard business review*, 84(10):92, 2006.
- [215] Michael Luca. Designing Online Marketplaces: Trust and Reputation Mechanisms. *Innovation Policy and the Economy*, 17:77–93, January 2017.
- [216] Troy J. Strader and Sridhar N. Ramaswami. The value of seller trustworthiness in C2C online markets. *Communications of the ACM*, 45(12):45–49, December 2002.
- [217] Luís Cabral and Ali Hortaçsu. The dynamics of seller reputation: Evidence from Ebay. *The Journal of Industrial Economics*, 58(1):54–78, March 2010.
- [218] Steven Tadelis. Reputation and Feedback Systems in Online Platform Markets. *Annual Review of Economics*, 8(1):321–340, October 2016.
- [219] Yuxin Chen and Song Yao. Sequential Search with Refinement: Model and Application with Click-Stream Data. *Management Science*, 63(12):4345–4365, December 2017.
- [220] Raluca M. Ursu. The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions. *Marketing Science*, 37(4):530–552, August 2018.
- [221] Babur De los Santos and Sergei Koulayev. Optimizing Click-Through in Online Rankings with Endogenous Search Refinement. *Marketing Science*, 36(4):542–564, July 2017.
- [222] Sergei Koulayev. Search for differentiated products: Identification and estimation. *The RAND Journal of Economics*, 45(3):553–575, September 2014.
- [223] Bo Xu, Zhangxi Lin, and Bingjia Shao. Factors affecting consumer behaviors in online buy-it-now auctions. *Internet Research*, 20(5):509–526, October 2010.
- [224] Daniel Houser and John Wooders. Reputation in Auctions: Theory, and Evidence from eBay. *Journal of Economics & Management Strategy*, 15(2):353–369, June 2006.
- [225] Aaron L Bodoh-Creed, Jörn Boehnke, and Brent Hickman. How efficient are decentralized auction platforms? *The Review of Economic Studies*, 88(1):91–125, 2021.

- [226] Patrick Bajari, , and Ali Hortacsu. The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions. *The RAND Journal of Economics*, 34(2):329, 22.
- [227] Lorenzo Dall’Amico, Antoine Fosset, Jean-Philippe Bouchaud, and Michael Benzaquen. How does latent liquidity get revealed in the limit order book? *Journal of Statistical Mechanics: Theory and Experiment*, 2019(1):013404, 2019.
- [228] Jean-Philippe Bouchaud, J. Doyne Farmer, and Fabrizio Lillo. Chapter 2 - how markets slowly digest changes in supply and demand. In Thorsten Hens and Klaus Reiner Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*, Handbooks in Finance, pages 57–160. North-Holland, San Diego, 2009.
- [229] Yuriy Gorodnichenko and Oleksandr Talavera. Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration. *American Economic Review*, 107(1):249–282, January 2017.
- [230] Aaron Bodoh-Creed, Jörn Boehnke, and Brent Hickman. Using Machine Learning to predict price dispersion. Working Paper, 2020.
- [231] David J. DiRusso, Susan M. Mudambi, and David Schuff. Determinants of prices in an online marketplace. *Journal of Product & Brand Management*, 20(5):420–428, August 2011.
- [232] Hee-Woong Kim and Yunjie Xu. Drivers of price Premium in e-markets. *Communications of the ACM*, 50(11):91–95, November 2007.
- [233] David Lucking-Reiley, Doug Bryan, Naghi Prasad, and Daniel Reeves. Pennies from Ebay: The determinants of price in online auctions. *Journal of Industrial Economics*, 55(2):223–233, June 2007.
- [234] Joe Van Buskirk, Amanda Roxburgh, Raimondo Bruno, Sundresan Naicker, Simon Lenton, Rachel Sutherland, Elizabeth Whittaker, Natasha Sindicich, Allison Matthews, Kerrynt Butler, and Lucinda Burns. Characterising dark net marketplace purchasers in a sample of regular psychostimulant users. *International Journal of Drug Policy*, 35:32–37, September 2016.
- [235] James Martin, Jack Cunliffe, David Décary-Hétu, and Judith Aldridge. Effect of restricting the legal supply of prescription opioids on buying through online illicit marketplaces: Interrupted time series analysis. *BMJ (Clinical research ed.)*, 361, 2018.
- [236] Sam Miller, Abeer El-Bahrawy, Martin Dittus, Mark Graham, and Joss Wright. Predicting Drug Demand with Wikipedia Views: Evidence from Darknet Markets. In *Proceedings of The Web Conference 2020*, pages 2669–2675, Taipei Taiwan, April 2020. ACM.
- [237] Judith Aldridge and David Décary-Hétu. Hidden wholesale: The drug diffusing ca-

- capacity of online drug cryptomarkets. *International Journal of Drug Policy*, 35:7–15, September 2016.
- [238] Marie Claire Van Hout and Tim Bingham. Responsible vendors, intelligent consumers: Silk Road, the online revolution in drug trading. *International Journal of Drug Policy*, 25(2):183–189, March 2014.
- [239] Márton Karsai, Nicola Perra, and Alessandro Vespignani. Time varying networks and the weakness of strong ties. *Scientific Reports*, 4(1):4001, May 2015.
- [240] Enrico Ubaldi, Nicola Perra, Márton Karsai, Alessandro Vezzani, Raffaella Burioni, and Alessandro Vespignani. Asymptotic theory of time-varying social networks with heterogeneous activity and tie allocation. *Scientific Reports*, 6(1):35724, December 2016.
- [241] A.L Barabási, H Jeong, Z Néda, E Ravasz, A Schubert, and T Vicsek. Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and its Applications*, 311(3-4):590–614, August 2002.
- [242] H Jeong, Z Néda, and A. L Barabási. Measuring preferential attachment in evolving networks. *Europhysics Letters (EPL)*, 61(4):567–572, February 2003.
- [243] European Monitoring Centre for Drugs and Drug Addiction. *Drugs and the Darknet: Perspectives for Enforcement, Research and Policy*. Publications Office, LU, 2017.
- [244] Gwern Branwen, Nicolas Christin, David Décary-Héту, Rasmus Munksgaard Andersen, StExo, El Presidente, Anonymous, Daryl Lau, Delyan Kratunov Sohlzl, Vince Cacic, Van Buskirk, Whom, Michael McKenna, and Sigi Goode. Dark net market archives, 2011-2015, July 2015.
- [245] European Monitoring Centre for Drugs and Drug Addiction. Reputation is everything: The role of ratings, feedback and reviews in cryptomarkets. In *The Internet and Drug Markets*. Publications Office, LU, 2016.
- [246] Robert Augustus Hardy and Julia R Norgaard. Reputation in the Internet black market: An empirical and theoretical analysis of the Deep Web. *Journal of Institutional Economics*, 12(3):515–539, 2016.
- [247] Scott W. Duxbury and Dana L. Haynie. Building them up, breaking them down: Topology, vendor selection patterns, and a digital drug market’s robustness to disruption. *Social Networks*, 52:238–250, January 2018.
- [248] Laura Alessandretti, Kaiyuan Sun, Andrea Baronchelli, and Nicola Perra. Random walks on activity-driven networks with attractiveness. *Physical Review E*, 95(5):052318, May 2017.
- [249] Austan D Goolsbee and Peter J Klenow. Internet rising, prices falling: Measuring inflation in a world of e-commerce. In *Aea papers and proceedings*, volume 108, pages

488–92, 2018.

- [250] Yuriy Gorodnichenko, Viacheslav Sheremirov, and Oleksandr Talavera. Price setting in online markets: Does it click? *Journal of the European Economic Association*, 16(6):1764–1811, 2018.
- [251] Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. Technical report, National Bureau of Economic Research, 2015.
- [252] Seth Siyuan Li and Elena Karahanna. Online recommendation systems in a B2C E-commerce context: A review and future directions. *Journal of the Association for Information Systems*, 16(2):2, 2015.
- [253] Antoine Mandel and Xavier Venel. Sequential competition and the strategic origins of preferential attachment. *International Journal of Game Theory*, pages 1–26, 2021.
- [254] Chih-Sheng Hsieh, Michael König, and Xiaodong Liu. Network formation with local complements and global substitutes: The case of r&d networks. *University of Zurich, Department of Economics, Working Paper*, (217), 2017.
- [255] Albert-László Barabási. The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039):207–211, May 2005.
- [256] Antoine Moinet, Michele Starnini, and Romualdo Pastor-Satorras. Burstiness and aging in social temporal networks. *Physical Review Letters*, 114(10):108701, March 2015.
- [257] Frank Wehinger. The Dark Net: Self-Regulation Dynamics of Illegal Online Markets for Identities and Related Services. In *2011 European Intelligence and Security Informatics Conference*, pages 209–213, Athens, Greece, September 2011. IEEE.
- [258] Coralio Ballester, Antoni Calvó-Armengol, and Yves Zenou. Who’s Who in Networks. Wanted: The Key Player. *Econometrica*, 74(5):1403–1417, September 2006.
- [259] William M Boal and Michael R Ransom. Monopsony in the labor market. *Journal of economic literature*, 35(1):86–112, 1997.
- [260] Arindrajit Dube, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. Monopsony in Online Labor Markets. *American Economic Review: Insights*, 2(1):33–46, March 2020.
- [261] Herbert London. The meaning of decentralization. *The Social Studies*, 66(2):55–59, 1975.
- [262] Nathan Schneider. Decentralization: an incomplete ambition. *Journal of Cultural Economy*, 12(4):265–285, 2019.
- [263] Serge Vaudenay. Centralized or decentralized? the contact tracing dilemma. Cryptology ePrint Archive, Paper 2020/531, 2020.

- [264] Matthieu Nadini, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. Mapping the nft revolution: market trends, trade networks, and visual features. *Scientific reports*, 11(1):1–11, 2021.
- [265] Abeer ElBahrawy, Laura Alessandretti, Anne Kandler, Romualdo Pastor-Satorras, and Andrea Baronchelli. Evolutionary dynamics of the cryptocurrency market. *Royal Society open science*, 4(11):170623, 2017.
- [266] Amin Mekacher, Alberto Bracci, Matthieu Nadini, Mauro Martino, Laura Alessandretti, Luca Maria Aiello, and Andrea Baronchelli. How rarity shapes the nft market. *arXiv preprint arXiv:2204.10243*, 2022.
- [267] Daniel Treisman. *The architecture of government: Rethinking political decentralization*. Cambridge University Press, 2007.
- [268] Mark Bray. Control of education: Issues and tensions in centralization and decentralization. *Comparative education: The dialectic of the global and the local*, 1999:207–232, 1999.
- [269] Robert D Ebel and Serdar Yilmaz. *On the measurement and impact of fiscal decentralization*, volume 2809. World Bank Publications, 2002.
- [270] Richard Saltman, Reinhard Busse, and Josep Figueras. *Decentralization in health care: strategies and outcomes*. McGraw-hill education (UK), 2006.
- [271] Alexander R Galloway. *Protocol: How control exists after decentralization*. MIT press, 2004.
- [272] Herbert G Tanner and Amit Kumar. Towards decentralization of multi-robot navigation functions. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 4132–4137. IEEE, 2005.
- [273] Linton C Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239, 1978.
- [274] Linton C. Freeman. Centrality in social networks conceptual clarification. *Social Networks*, 1(3):215–239, January 1978.
- [275] Governance, 2022. Accessed June 6, 2022.
- [276] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June Hsu, and Kuansan Wang. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*, pages 243–246, 2015.
- [277] Next steps for microsoft academic – expanding into new horizons, 2021. Accessed June 6, 2022.

- [278] Semantic scholar’s paper field of study classifier. [https://github.com/allenai/s2\\_fos](https://github.com/allenai/s2_fos) Accessed June 8, 2022, 2022.
- [279] Charles C. Hyland, Yuanming Tao, Lamiae Azizi, Martin Gerlach, Tiago P. Peixoto, and Eduardo G. Altmann. Multilayer networks for text analysis with multiple data types. *EPJ Data Science*, 10(1):33, December 2021.
- [280] Staša Milojević. Quantifying the cognitive extent of science. *Journal of Informetrics*, 9(4):962–973, October 2015.
- [281] Lutz Bornmann, Robin Haunschild, and Rüdiger Mutz. Growth rates of modern science: A latent piecewise growth curve approach to model publication numbers from established and new literature databases. *Humanities and Social Sciences Communications*, 8(1):224, December 2021.
- [282] J Toulmin Smith. Local self-government and centralization: The characteristics of each: And its practical tendencies as affecting social , moral, and political welfare and progress, including comprehensive outlines of the english constitution. *Moral, and Political Welfare and Progress (London, 1851)*, pages 27–51, 1851.
- [283] James M Buchanan. Individual choice in voting and the market. *Journal of Political Economy*, 62(4):334–343, 1954.
- [284] Harold F Kaufman. Toward an interactional conception of community. *Soc. F.*, 38:8, 1959.
- [285] Peter Bachrach and Morton S Baratz. Two faces of power1. *The American Political Science Review*, 56(4):947–952, 1962.
- [286] D.D. Šiljak. On decentralized control of large scale systems. *IFAC Proceedings Volumes*, 11(1):1849–1856, 1978. 7th Triennial World Congress of the IFAC on A Link Between Science and Applications of Automatic Control, Helsinki, Finland, 12-16 June.
- [287] Louis Columbus. Gartner hype cycle for emerging technologies, 2016 adds blockchain & machine learning for first time, 2016. Accessed June 9, 2022.
- [288] Sue McLean and Simon Deane-Johns. Demystifying blockchain and distributed ledger technology—hype or hero? *Computer Law Review International*, 17(4):97–102, 2016.
- [289] Ferdinando M Ametrano. Bitcoin, blockchain, and distributed ledgers: Between hype and reality. *Available at SSRN 2832249*, 2016.
- [290] Alberto Bracci<sup>+</sup>, Matthieu Nadini<sup>+</sup>, Maxwell Aliapoulios, Damon McCoy, Ian Gray, Alexander Teytelboym, Angela Gallo, and Andrea Baronchelli. Dark Web Marketplaces and COVID-19: After the vaccines. *arXiv preprint arXiv:2102.05470v4*, 2021. (+) Contributed equally.

- [291] Chris Isidore. Feds seize 1 billion in Bitcoins they say were stolen from silk road. <https://edition.cnn.com/2020/11/06/business/bitcoin-seized-silk-road-ulbricht/index.html> Accessed January 4, 2021, 2020. CNN.
- [292] Chainalysis. The Chainalysis 2021 crypto crime report. <https://go.chainalysis.com/2021-Crypto-Crime-Report.html> Accessed October 27, 2021, 2021.
- [293] Naoki Hiramoto and Yoichi Tsuchiya. Measuring dark web marketplaces via bitcoin transactions: From birth to independence. *Forensic Science International: Digital Investigation*, 35:301086, 2020.
- [294] Julia Buxton and Tim Bingham. The rise and challenge of dark net drug markets. *Policy brief*, 7:1–24, 2015.
- [295] Alexia Maddox, Monica J Barratt, Matthew Allen, and Simon Lenton. Constructive activism in the dark web: cryptomarkets and illicit drugs in the digital ‘demimonde’. *Information, Communication & Society*, 19(1):111–126, 2016.
- [296] Atte Oksanen, Bryan Lee Miller, Iina Savolainen, Anu Sirola, Jakob Demant, Markus Kaakinen, and Izabela Zych. Illicit drug purchases via social media among american young people. In *International Conference on Human-Computer Interaction*, pages 278–288. Springer, 2020.
- [297] German police seized nine telegram-based drug markets. <https://darknetlive.com/post/german-police-seized-nine-telegram-based-drug-markets/> Accessed October 27, 2021, 2020. DarknetLive.
- [298] Yik-Hei Sung, Wing-Ho Lee, Franco Ka-Wah Leung, and Jonathan J Fong. Prevalence of illegal turtle trade on social media and implications for wildlife trade monitoring. *Biological Conservation*, 261:109245, 2021.
- [299] Andrew Childs, Melissa Bull, and Ross Coomber. Beyond the dark web: Navigating the risks of cannabis supply over the surface web. *Drugs: Education, Prevention and Policy*, pages 1–12, 2021.
- [300] K Hazel Kwon and Chun Shao. Dark knowledge and platform governance: A case of an illicit e-commerce community in reddit. *American Behavioral Scientist*, 65(6):779–799, 2021.
- [301] Rasmus Munksgaard, James Martin, et al. How and why vendors sell on cryptomarkets. *Trends and Issues in Crime and Criminal Justice*, 608:1, 2020.
- [302] Silje Anderdal Bakken and Jakob Johan Demant. Sellers’ risk perceptions in public and private social media drug markets. *International Journal of Drug Policy*, 73:255–262, 2019.

- [303] European Monitoring Centre for Drugs, Drug Addiction, and Europol. Drugs and the darknet: Perspectives for enforcement, research and policy, 2017.
- [304] World Health Organization. World drug report 2019. *United Nations publication, Sales No. E, 19*, 2019.
- [305] Gwern. Updated: list of dark net markets (Tor & I2P). <https://www.gwern.net/docs/sr/2019-04-22-deepdotweb-dnmlist.html> Accessed October 27, 2021, 2020.
- [306] Aaron W Baur, Julian Bühler, Markus Bick, and Charlotte S Bonorden. Cryptocurrencies as a disruption? Empirical findings on user adoption and future potential of Bitcoin and co. In *Conference on e-Business, e-Services and e-Society*, pages 63–80. Springer, 2015.
- [307] Ed Saiedi, Anders Broström, and Felipe Ruiz. Global drivers of cryptocurrency infrastructure adoption. *Small Business Economics*, pages 1–54, 2020.
- [308] Joe Van Buskirk, Sundresan Naicker, RB Bruno, C Breen, and A Roxburgh. Drugs and the internet. *The National Illicit Drug Indicators Project*, 2016.
- [309] Frank Wilcoxon. Individual Comparisons by Ranking Methods. In *Breakthroughs in Statistics*, pages 196–202. Springer, 1992.
- [310] Matthieu Nadini. Temporalbackbone. <https://pypi.org/project/TemporalBackbone/> Accessed October 27, 2021, 2021. Python pip 3 library: A tool to detect the backbone in temporal networks.
- [311] Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, pages 50–60, 1947.
- [312] Frank J Massey Jr. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American statistical Association*, 46(253):68–78, 1951.
- [313] Charles Spearman. The proof and measurement of association between two things. *Appleton-Century-Crofts*, 1961.
- [314] Andréanne Bergeron, David Décary-Héту, and Luca Giommoni. Preliminary findings of the impact of COVID-19 on drugs crypto markets. *International Journal of Drug Policy*, 83:102870, 2020.
- [315] Covid is causing shipping issues, but natural competitive forces are causing darknet market consolidation. <https://blog.chainalysis.com/reports/darknet-markets-cryptocurrency-2020> Accessed October 27, 2021, 2020. Chainalysis Team.
- [316] Abhineet Gupta, Sean B Maynard, and Atif Ahmad. The dark web phenomenon: A review and research agenda. *arXiv preprint arXiv:2104.07138*, 2021.

- [317] Martin Horton-Eddison, Patrick Shortis, Judith Aldridge, and Fernando Caudevilla. Drug cryptomarkets in the 2020s: Policy, enforcement, harm, and resilience. *Global Drug Policy Observatory*, 2021.
- [318] Office of Public Affairs. First nationwide undercover operation targeting darknet vendors results in arrests of more than 35 individuals selling illicit goods and the seizure of weapons, drugs and more than \$23.6 million. <https://www.justice.gov/opa/pr/first-nationwide-undercover-operation-targeting-darknet-vendors-results-arrests-mo> Accessed October 27, 2021, 2018. Department of Justice, United States.
- [319] Europol Team. International sting against dark web vendors leads to 179 arrests. <https://www.europol.europa.eu/newsroom/news/international-sting-against-dark-web-vendors-leads-to-179-arrests> Accessed October 27, 2021, 2020. Europol.
- [320] Alberto Bracci, Jörn Boehnke, Abeer ElBahrawy, Nicola Perra, Alexander Teytelboym, and Andrea Baronchelli. Macroscopic properties of buyer-seller networks in online marketplaces. *arXiv preprint arXiv:2112.09065*, 2021.
- [321] Dune analytics. <https://dune.xyz/rchen8/opensea> Accessed June 22, 2022, 2022.
- [322] NFT beats cheugy to be Collins Dictionary’s word of the year. <https://www.theguardian.com/books/2021/nov/24/nft-is-collins-dictionary-word-of-the-year> Accessed March 15, 2022, 2021.
- [323] How to create an nft collection with a smart contract. <https://geekflare.com/finance/create-nft-collection-with-a-smart-contract/> Accessed April 19, 2022, 2022.
- [324] Bored apes and cryptopunks help jolt nft market to over 21,000% growth and \$17.6 billion in sales last year. <https://fortune.com/2022/03/10/bored-apes-cryptopunks-jolt-nft-market-to-billions-in-sales/> Accessed March 27, 2022, 2022.
- [325] Generative art. <https://www.tate.org.uk/art/art-terms/g/generative-art> Accessed March 28, 2022, 2022.
- [326] CryptoPunks. <https://www.larvalabs.com/cryptopunks> Accessed April 1, 2022, 2017.
- [327] Bored Ape Yacht Club. <https://boredapeyachtclub.com/#/> Accessed 25 March, 2022, 2021.
- [328] rarity.tools. <https://rarity.tools> Accessed June 30, 2022.
- [329] Highest-ever CryptoPunk sale nets \$24 million. <https://www.theblockcrypto.com/linked/134139/highest-ever-cryptopunk-sale-nets-24-million> Accessed

- March 19, 2022, 2022.
- [330] Luisa Schaar and Stylianos Kampakis. Non-fungible tokens as an alternative investment: Evidence from cryptopunks. *The Journal of The British Blockchain Association*, page 31949, 2022.
- [331] Rare Bored Ape Yacht Club NFT Sells for Record \$3.4 Million USD. <https://hypebeast.com/2021/10/bored-ape-yacht-club-nft-3-4-million-record-sothebys-metaverse> Accessed March 19, 2022, 2021.
- [332] Kenneth Koford and Adrian E Tschoegl. The market value of rarity. *Journal of Economic Behavior & Organization*, 34(3):445–457, 1998.
- [333] Yeonjoon Lee. Measuring the impact of rarity on price: Evidence from nba top shot. *Available at SSRN 3906231*, 2021.
- [334] Jonathan E Hughes. Demand for rarity: Evidence from a collectible good. *downloaded from, (Year: 2021)*, 2020.
- [335] Soroush Ghazi and Mark Schneider. Preference for scarcity. *Available at SSRN*, 2021.
- [336] De-Rong Kong and Tse-Chun Lin. Alternative investments in the fintech era: The risk and return of non-fungible token (nft). *Available at SSRN 3914085*, 2021.
- [337] OpenSea. <https://opensea.io/> Accessed July 3, 2022.
- [338] OpenSea Assets. <https://opensea.io/assets> Accessed April 3, 2022.
- [339] Foundation. <https://foundation.app/> Accessed 25 March, 2022.
- [340] SuperRare. <https://superrare.com/> Accessed 25 March, 2022.
- [341] Nifty gateway. <https://niftygateway.com/> Accessed April 2, 2022.
- [342] Feral file. <https://feralfile.com/exhibitions> Accessed April 2, 2022.
- [343] Buying nfts during presales and public mints: Things you should know. <https://www.coindesk.com/learn/buying-nfts-during-presales-and-public-mints-things-you-should-know/> Accessed April 15, 2022, 2022.
- [344] Nft collections listed by sales volume | coinmarketcap. <https://coinmarketcap.com/nft/collections/> Accessed June 18, 2022.
- [345] rarity.tools. Ranking Rarity: Understanding Rarity Calculation Methods. <https://raritytools.medium.com/ranking-rarity-understanding-rarity-calculation-methods-86ceaeb9b98c> Accessed March 20, 2022, 2021.

- [346] How azukis suddenly became the world's best-selling nft collection. <https://www.forbes.com/sites/jeffkaufman/2022/02/14/how-azukis-suddenly-became-the-worlds-best-selling-nft-collection/?sh=5e6aff5a5cbe> Accessed April 20, 2022, 2022.
- [347] The original cryptopunk. <https://andrewschrbr.medium.com/the-original-cryptopunk-a852101a03ed> Accessed April 4, 2022, 2021.
- [348] Deborah Rabinowitz. Seven forms of rarity. *Biological aspects of rare plant conservation*, 1981.
- [349] Cyrille Violle, Wilfried Thuiller, Nicolas Mouquet, François Munoz, Nathan JB Kraft, Marc W Cadotte, Stuart W Livingstone, and David Mouillot. Functional rarity: the ecology of outliers. *Trends in Ecology & Evolution*, 32(5):356–367, 2017.
- [350] Duckling open-source library. <https://github.com/facebook/duckling> Accessed January 4, 2021, 2018.
- [351] Cryptocompare api. <https://min-api.cryptocompare.com/> Accessed January 4, 2021, 2020.
- [352] Foreign exchange rates api with currency conversion. <https://exchangeratesapi.io> Accessed January 4, 2021, 2019.
- [353] Jeffrey D Scargle, Jay P Norris, Brad Jackson, and James Chiang. Studies in astronomical time series analysis. VI. Bayesian block representations. *The Astrophysical Journal*, 764(2):167, 2013.
- [354] Katharina Zeuch. An introduction to generative art nfts. <https://medium.com/@datash/an-introduction-to-generative-art-nfts-35e650a0f281> Accessed March 15, 2022.
- [355] Inside the punk art machinery - how to generate 10 000 punks (and punkettes), algorithmically - paint by numbers. [https://github.com/cryptopunksnotdead/programming-cryptopunks/blob/master/03\\_generate.md](https://github.com/cryptopunksnotdead/programming-cryptopunks/blob/master/03_generate.md) Accessed April 4, 2022, 2021.
- [356] Creating a 10k nft avatar collection: Expansion-punks step-by-step. <https://medium.com/geekculture/expansionpunks-welcome-to-the-more-inclusive-punkverse-aa77e675bcf5> Accessed April 4, 2022, 2021.