Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-up, Stable, and Covid-19 Eras

Anh V. Vu, Jack Hughes, Ildiko Pete, Ben Collier, Yi Ting Chua, Ilia Shumailov, Alice Hutchings
firstname.lastname@cl.cam.ac.uk
Cambridge Cybercrime Centre, Department of Computer Science & Technology
University of Cambridge, Cambridge, CB3 0FD, UK

ABSTRACT

Trust and reputation play a core role in underground cybercrime markets, where participants are anonymous and there is little legal recourse for dispute arbitration. These underground markets exist in tension between two opposing forces: the drive to hide incriminating information, and the trust and stability benefits that greater openness yields. Revealing information about transactions to mitigate scams also provides valuable data about the market. We analyse the first dataset, of which we are aware, about the transactions created and completed on a well-known and high-traffic underground marketplace, Hack Forums, along with the associated threads and posts made by its users over two recent years, from June 2018 to June 2020. We use statistical modelling approaches to analyse the economic and social characteristics of the market over three eras, especially its performance as an infrastructure for trust. In the Set-up era, we observe the growth of users making only one transaction, as well as 'power-users' who make many transactions. In the Stable era, we observe a wide range of activities (including large-scale transfers of intermediate currencies such as Amazon Giftcards) which declines slowly from an initial peak. Finally, we analyse the effects of the Covid-19 pandemic, concluding that while we see a significant increase in transactions across all categories, this reflects a *stimulus* of the market, rather than a *transformation.* New users overcome the 'cold start' problem by engaging in low-level currency exchanges to prove their trustworthiness. We observe *currency exchange* accounts for most contracts, and Bitcoin and PayPal are the preferred payment methods by trading values and number of contracts involved. The market is becoming more centralised over time around influential users and threads, with significant changes observed during the Set-up and Covid-19 eras.

CCS CONCEPTS

• **Social and professional topics** → **Computer crime**; • **Mathematics of computing** → Time series analysis; • **Security and privacy** → Social aspects of security and privacy.

KEYWORDS

underground economy; hacking forums; cybercrime market; economic evolution; coronavirus; Covid-19; pandemic

IMC '20, October 27–29, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8138-3/20/10.
https://doi.org/10.1145/3419394.3423636

ACM Reference Format:
Anh V. Vu, Jack Hughes, Ildiko Pete, Ben Collier, Yi Ting Chua, Ilia Shumailov, Alice Hutchings. 2020. Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-up, Stable, and Covid-19 Eras. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3419394.3423636

1 INTRODUCTION

Online illicit marketplaces are a key part of the cybercrime economy, enabling malicious actors to cash out earnings, trade in malware, and obtain compromised credit cards. Trust and reputation are key aspects of any market, but particularly so in cybercrime markets, where user anonymity means that 'ripping' or defrauding customers is common. Underground markets have adapted to this problem of 'lemonisation' by providing reputation and vouching systems, escrow systems, and verified status to minimise information asymmetry [10, 11, 16, 17, 21, 27]. While previous research has focused on vendor ratings and public feedback, there are few datasets as comprehensive as the one we analyse in this paper.

On the longest-running and most popular cybercrime forum, a market has been active for some time. Previously, users posted advertisements while transactions were finalised off-site. Hack Forums does not officially offer a formal escrow service, however, in reaction to widespread concerns about abuse, the administrators recently (June 2018) opened a dedicated marketplace to facilitate the exchange and trade in goods and services where contracts are logged. This appears to be primarily used as a reputation and trust management system, in which transaction details are visible to forum users on payment of a small fee. This trust adaptation presents a unique opportunity for academic research. In this paper, we present an extensive analysis of this burgeoning market. We are particularly interested in exploring the longitudinal evolution of conflict, trust, and activity of different kinds – we consider this marketplace to be an example of a disparate group of actors coming together (with their own diverse motivations) to work on a joint endeavour. Thus, we draw on Tuckman's stages of group development (discussed in Section 2.2) to guide our analysis.

Our dataset provides valuable insights into the economic activity linked to the forum, and how an underground marketplace evolves. The dataset begins from the start of this system, with details of goods traded, currency types and amount, time taken for transactions to complete, and in some cases Bitcoin addresses for transactions. First, we describe the dataset and its characteristics in §3. As this kind of 'social' data is not originally generated for research purposes, it requires additional interpretation, so we include findings from our exploratory analysis of the data in this section where they are useful for making sense of the dataset and what it represents. We then provide a descriptive economic analysis, outlining how much money is being made, by how many users, and from what goods and services (§4). We also identify the preferred payment methods used and the top merchandise being traded. In §5, we use statistical approaches to explore the market's longitudinal evolution, how users overcome the 'cold start' problem, as well as the trust and reputation functions facilitated by this marketplace. In §6, we draw major insights observed from our analyses over three eras. We then discuss the contributions to the security community and the limitations of our work in §7. Ethical issues are discussed in the Appendix. The dataset is made available for academia through data sharing agreements so that other academic researchers can fully reproduce our experiments and build further analyses based on our results.

2 ANALYSING ONLINE MARKETPLACES

The products and services available in online underground marketplaces are diverse, ranging from illicit drugs [6], malicious software [12], to stolen data [7, 10, 11]. Forums remain a popular platform as they provide an easy way to establish business and social connections [27]. As platforms, forums contain features such as reputation systems and hierarchical levels of moderators and administrators that allow for the exertion of social control on the community [1, 7, 16, 27]. A pivotal question is how members of anonymous illicit marketplaces can trust each other. This is important to the sustainability of a marketplace, as trust increases actual purchases [8, 20]. Underground markets use various techniques to facilitate trust, including vendor verification and reputation systems [10, 16, 27]. Moderators promote trust by verifying users and excluding scammers [7, 11, 16]. Member activity may also act as a signal of trust, as experienced users tend to be perceived as trustworthy [7, 10, 11, 27]. While existing literature mainly relied on feedback from users, our research provides insights on the topic with contractual data on created and completed transactions.

2.1 Related Work

Machine learning and natural language processing (NLP) methods have been used to analyse marketplaces at scale. For example, Sun et al. [23] propose a machine learning-based approach to detect private interactions on the Nulled forum, where they examine the trading activities and monetisation methods of members. NLP methods have been used to identify the function and intent of messages [5], identify posts related to transactions and to extract products and prices [22], and to identify supply chains [3].

Afroz et al. [1] identify common features across successful forums, including top-down governance, norm-conforming behaviors from members, frequent communication, and the use of enforcement such as fines and bans. Holt [9] qualitatively analysed ten publicly accessible Russian forums, finding the relationship between members is influenced by price, customer service, and trust. Allodi et al. [2] investigated factors contributing to the success of online markets by analysing marketplace models, preferences of online traders, market stability, and resiliency of cybercrime tools.

A subset of research takes a network-based approach in analysing the dynamics of transactions within underground marketplaces. The study by Motoyama et al. [16] analyses six underground forums (not including Hack Forums) by building the social networks based on private messages, friend status and thread posting behaviour. Their dataset contains descriptive information, such as forum posts, private messages, user logs and user registration data; however, it lacks transactions made by the forum members.

2.2 Theoretical Approach

To examine the evolution and maturation of the Hack Forums marketplace, we split the timespan into three eras. We draw on Tuckman's [24] stages of group development, which proposes that established groups go through four stages: *forming*, *storming*, *norming*, and *performing*. Our first two 'eras' are Set-up (*forming* and *storming*), from when the contract system was adopted (1 June 2018) to before the time contracts become mandatory (1 March 2019), and Stable (*norming*), from the end of Set-up to before 11 March 2020. At the beginning of 2020, Covid-19 began to spread globally, with a global pandemic declared by the World Health Organisation (WHO) on 11 March 2020 [26]. This coincides with an uptick in the number of contracts on the forum, representing the *performing* stage. We name this most recent era Covid-19, which spans from 11 March 2020 to the end of the data collection period on 30 June 2020. In some figures, we denote the three eras as E1, E2, and E3.

It bears noting that the events which we use to define these eras play a *deductive* rather than *inductive* role in our analysis – they are imposed by us in order to analyse the effects of external factors, rather than arising from the data. We use Tuckman's stages in our analysis to make sense of the changes we observe in our data which correspond to these external events; namely, how they appear to reflect the evolution of trust, collaboration and conflict in this marketplace. Tuckman's stages of group development assumes group membership is static. However, as with many online networks, users of underground markets are transient. Therefore, we also consider how users overcome the *cold start problem* [13, 15] in this context. Here, the 'cold start' problem refers to the difficulties faced by new users who find that others do not want to trade with them due to lack of reputation, but cannot gain reputation as nobody will trade with them. While the 'cold start' problem is a well-recognised issue for recommendation systems, we believe we are the first to consider it in relation to underground markets.

3 DATASET

We provide the first analysis, of which we are aware, of contractual transactions made in the underlying marketplace on one of the most high-traffic and well-known online underground forums Hack Forums. Our data, newly collected as a part of the CrimeBB dataset [19], contains nearly 190,000 real contracts created by users over two years from June 2018 to June 2020. The contracts represent transactions, and some are linked with the advertising threads and discussion posts which provides additional context. Each contract includes the goods and services being exchanged, obligations, agreement terms, and the ratings of the parties involved.

The Contract System. On Hack Forums, a contract is an agreement between members for trading goods or services. Contracts were optional when first introduced to the forum in June 2018. However, on 30 January 2019, it was announced that contracts would be *mandatory* for all deals from 1 March 2019. While some

Table 1: Taxonomy of collected contracts from June 2018 to June 2020

Type\Status	Complete	Active Deal	Disputed	Incomplete	Cancelled	Denied	Expired	Total
Sale	39,908 (21.20%)	1,931 (1.03%)	1,009 (0.54%)	66,347 (35.25%)	6,795 (3.61%)	64 (0.03%)	6,080 (3.23%)	122,134 (64.88%)
Purchase	11,893 (6.32%)	10 (0.01%)	629 (0.33%)	4,703 (2.50%)	2,378 (1.26%)	29 (0.02%)	2,761 (1.47%)	22,403 (11.90%)
Exchange	28,157 (14.96%)	2 (0.00%)	455 (0.24%)	3,342 (1.78%)	5,758 (3.06%)	66 (0.04%)	2,588 (1.37%)	40,368 (21.45%)
Trade	1,325 (0.70%)	1 (0.00%)	21 (0.01%)	547 (0.29%)	197 (0.10%)	3 (0.00%)	256 (0.14%)	2,350 (1.25%)
Vouch Copy	566 (0.30%)	0 (0.00%)	3 (0.00%)	228 (0.12%)	56 (0.03%)	0 (0.00%)	128 (0.07%)	981 (0.52%)
Total	81,849 (43.48%)	1,944 (1.03%)	2,117 (1.12%)	75,167 (39.93%)	15,184 (8.07%)	162 (0.09%)	11,813 (6.28%)	188,236 (100%)

Table 2: Visibility of contract types

Type\Visibility	Private	Public	Total
Sale Created	112,377 (92.01%)	9,757 (7.99%)	122,134
Purchase Created	17,723 (79.11%)	4,680 (20.89%)	22,403
Exchange Created	33,064 (81.91%)	7,304 (18.09%)	40,368
Trade Created	1,741 (74.09%)	609 (25.91%)	2,350
Vouch Copy Created	798 (81.35%)	183 (18.65%)	981
Sale Completed	35,099 (87.95%)	4,809 (12.05%)	39,908
Purchase Completed	9,013 (75.78%)	2,880 (24.22%)	11,893
Exchange Completed	23,461 (83.32%)	4,696 (16.68%)	28,157
Trade Completed	974 (73.51%)	351 (26.49%)	1,325
Vouch Copy Completed	466 (82.33%)	100 (17.67%)	566

transactions may be completed outside the contract system, in private channels such as direct messages, the regulation was strictly adopted, as it was announced that those avoiding the system would face account closure. This avoidance, if it occurs, might lead to inaccurate measurements in this study, however, as the system allows private contracts which do not reveal the goods being exchanged, we believe users are incentivised to use the system as it enables them to gain reputation and provides a certain level of protection (e.g., opening disputes) if anything goes wrong with a transaction.

Contract Process. To create a contract, the *maker* specifies the details and the user they want to trade with. If the receiving party denies the proposed contract, it becomes *denied*. If they accept, the contract becomes an *active deal* and they become the *taker*. Otherwise, the contract is marked as *expired* after 72 hours if no decision is made. After both parties accept the contract's terms and obligations and complete their own obligations, they can mark the contract as *complete* and users can rate each other (commonly known as *B-rating*). If either party is unsatisfied with the deal, they can open a dispute. The detailed process is shown in the Appendix.

Contract Taxonomy. We observe all contracts belong to one of five types: Sale, Purchase, and Vouch Copy are one-way, while Exchange and Trade are bi-directional. For the economic analysis, we exclude Vouch Copy, which was recently introduced in February 2020, as it represents a proof of reputation, not an economic trade. Table 1 shows the number and proportion of contracts for each type among *all* collected data. Overall, Sale dominates the others, accounting for 64.9% of contracts created, around three times higher than Exchange (21.5%), but has the highest non-completion rate (54.3%). Purchase, the reverse type of Sale, accounts for 11.9% of contracts, while Trade accounts for only 1.3%. Exchange has the highest completion rate, at 69.8%, more than double the completion rate of Sale (32.7%), indicating that Exchange are more likely to be accepted and settled. Vouch Copy is the only type with no denials.

Contract Visibility. Contracts can be public (users with an upgraded account can view all details) or private (some information is restricted to involved parties). Information available relating to private contracts include the maker, taker, type of deal, created date, and expiry date. Public contracts also include the obligations of each party, terms, goods to be exchanged, and ratings. If a user opens a dispute, the contract becomes public regardless of its previous visibility. Table 2 shows the visibility of contracts by category. Among created contracts, the proportion of public and private contracts is 12.0% and 88.0% respectively. For completed transactions, the percentage of public contracts is about 30% higher, accounting for 15.7% of contracts, with 84.3% remaining private. This suggests users tend to hide the majority of their contract details. Public contracts are more likely to be settled, with 57.0% transactions completed compared to 41.7% in private contracts. For both created and completed transactions, while the proportions of public Purchase, Exchange, Trade, Vouch Copy over the total are around 20%, the percentage of public Sale among all Sale contracts is considerably smaller, accounting for 8.0% and 12.1%, respectively. This indicates that contracts created by the sellers are more likely to be private.

Threads and Posts. To advertise goods or services, traders often create a thread describing their offerings, which can then be associated with a contract. Not all of the threads associated with contracts are advertisements, some are more general discussion threads from elsewhere on the forum. In our dataset, we observe 68.4% of *public* contracts (8.2% overall) are associated with a thread. Our dataset includes around 6,000 threads containing roughly 200,000 posts made by nearly 30,000 members from June 2018 to June 2020.

4 THE UNDERGROUND ECONOMY

In this section, we describe the evolution of the Hack Forums marketplace in terms of trading activities, payment methods, transaction values, number of contracts, parties involved, and completion time of contracts. We also examine market centralisation over time by looking at the social network formed by contractual relationships between users. Note that for analyses relying on contractual obligations, we only use *public* contracts, as the information is hidden in *private* ones.

4.1 Members and Contracts

Figure 1 shows an unstable fluctuation in the monthly growth of new contracts and new members who are party to a contract over the three eras. Overall, while there are significant shifts between

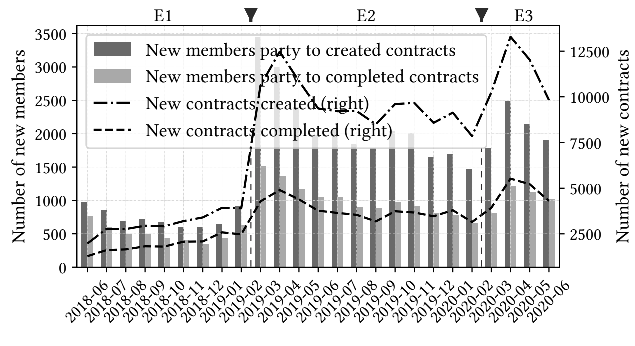

Figure 1: Monthly growth of new members and contracts

the eras, the number of new contracts created and new members tend to fluctuate together, except during the Set-up era, when the number of new contracts gradually grew but the number of new members joining the marketplace moderately decreased. During the 9-month Set-up era, the number of monthly created and completed contracts roughly doubled, despite the number of new members joining gradually decreasing. This indicates that on average, users were making more contracts each month during this era.

The Stable era begins with a policy change requiring contracts for all marketplace transactions. Compared to the month before, created contracts increase by 172% and completed contracts by 73%. A peak in April 2019, likely due to adoption of the new regulation, with around 12,500 contracts made and 5,000 completed, is followed by a gradual decline, to around 8,000 contracts made and 3,000 completed per month. From the end of Set-up to the end of Stable, the number of monthly created contracts doubles, but completed contracts only increase by 27%. At the beginning of this era, many new members start participating, peaking in March 2019 with 276% and 143% more new members involved in creating and completing contracts compared to the month before. The participation of new members then moderately declines to less than 50% of the peak at the end of Stable with around 1,500 and 700 users, respectively.

The last 4-months of data collection in the Covid-19 era show a sharp but fairly short-lived peak in both new contracts and new members joining the market. In April 2020, there are more than 13,000 created and 5,500 completed contracts, even surpassing the peak during the Stable era. While the number of new members also increases, this does not outpace the past peak, indicating established members are contributing more at this time. During this era, the ratio of the new members involved in created and completed contracts stays unchanged, with nearly 50% users involved in contracts not completing. This indicates while there is a stimulus in the Covid-19 era, the users involvement of the market remains stable. After the peak in April 2020, we see a drop in both number of users and contracts, showing a decrease of trading activities on the marketplace. It appears the lockdown intensively affected to the market for only a short period after the pandemic was declared.

Contract Visibility. Figure 2 shows the proportion of created and completed *public* contracts declines over the three eras. The proportion of completed public contracts is consistently higher, indicating public contracts are more likely to be completed. The biggest shift is in the Set-up era, when the percentage of public contracts began

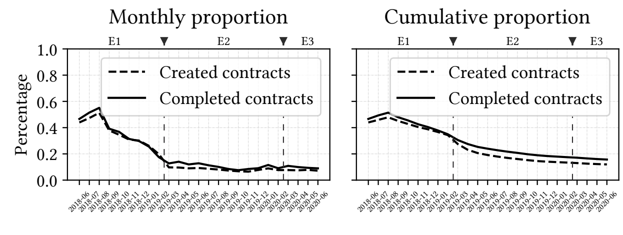

Figure 2: The proportion of *public* contract over time

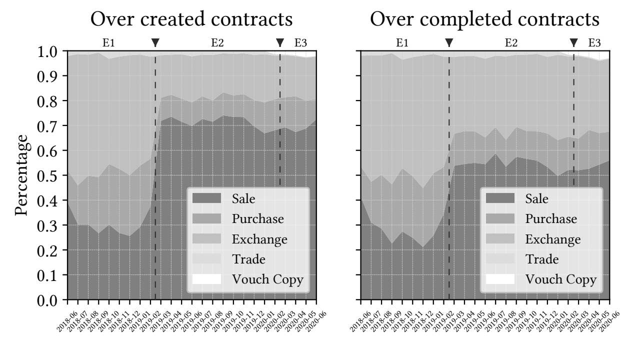

Figure 3: Contract type proportions by months

at around 45%, peaking in August 2018 at over 50%, then decreasing to around 20%. At the beginning of Stable, when contracts are made compulsory, the proportion dropped again, accounting for around 10% then remained mostly unchanged afterwards.

Contract Types. Figure 3 shows the evolution of monthly proportion of contract types from June 2018 to June 2020. For both created and completed contracts, the market is mostly occupied by Sale, Purchase, and Exchange. Trade and Vouch Copy constantly account for a small proportion (mostly less than 2%). During Set-up, the proportion of created and completed contract types stays mostly the same. At the beginning, Exchange accounts for the largest proportion (around 50%), followed by Sale (about 40%). Purchase starts around 10%, then gradually increases over time.

At the beginning of the Stable era, the market composition shifts, with Sale and Exchange swapping positions. Sale dominates the other types, accounting for over 70% of created and 55% of completed contracts. The percentage of Exchange declines to less than 20% of created and 30% of completed contracts. The proportion of Purchase also drops to around 10% and 15%, respectively. In this era, although the ordering of contract types stays the same, the proportion of completed Sale is lower than completed Exchange, indicating Exchange is more likely to be completed.

Despite the increase in contracts and members in Covid-19, we observe little change in the proportion of contract types, suggesting a market *stimulus* rather than a *transformation.* At the end of this era, Sale still dominate, accounting for over 70% of created, and 55% of completed contracts. Vouch Copy, adopted in February 2020, rapidly outpaces Trade, and continues to increase. In June 2020, the number of Vouch Copy increased by around 91% for created and 160% for completed contracts, compared to the end of Stable (February 2020). This rise suggests an increased desire to establish reputation within the market, which has a known history of scams.

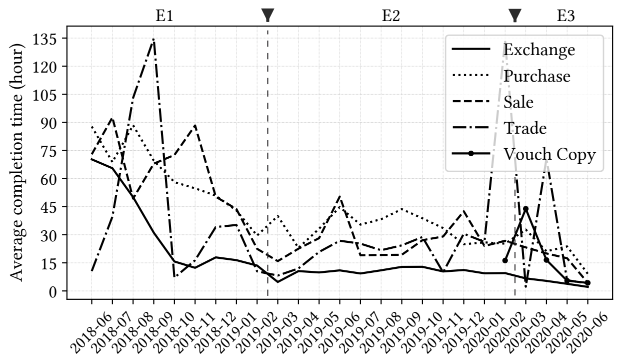

Figure 4: Average completion time by contract types

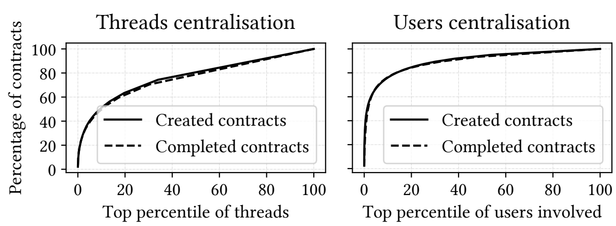

Figure 5: Top percentile of threads and users involved

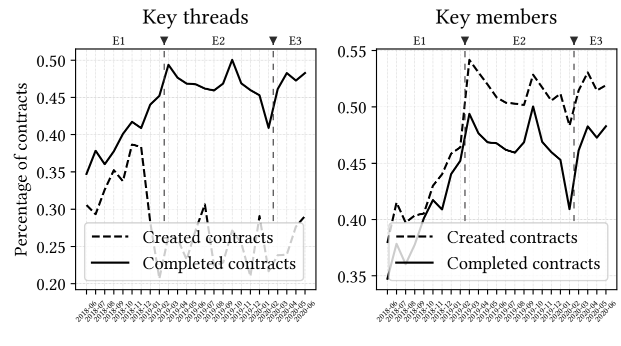

Figure 6: Key thread/member proportion by months

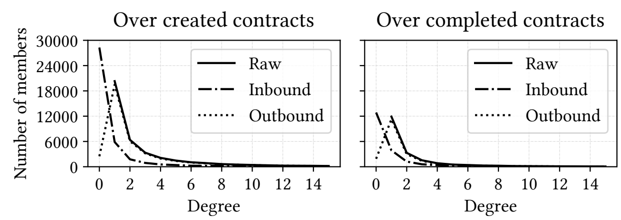

Figure 7: Degree distribution of the contractual network

Contract Completion Time. Overall, contracts complete faster over time (see Figure 4). The maximum completion times occur in the early months of Set-up, with the exception of Vouch Copy, which also saw a drop in completion time after its introduction in February 2020. The minimum completion times for all contract types occur during Covid-19, with contracts taking less than 10 hours to complete in June 2020. This increased speed over time is presumably due to users becoming more familiar with the contract system, or to increased trading demand. There is a drop from Set-up to Stable for all types; likewise, during Covid-19, completion times for most contract types decrease, except for abnormal short-lived peaks observed in Trade in February and April 2020. Along with quickly completed contracts, we see some very time-consuming transactions of Trade. As the number of Trade is rather small, it is likely that those peaks are due to noise. We note for this analysis we only consider contracts in which the completion date is provided, however, as these contracts account for around 70% of all completed contracts, we believe this is representative of the marketplace.

4.2 Market Centralisation

In most communities, centralisation naturally occurs around highly influential members. Our dataset shows that the market is largely centralised, with a small number of threads and members covering the majority of transactions. Figure 5 shows the relationship between the percentage of contracts made and the corresponding top percentile of threads and members involved. For both created and completed contracts, about 5% of users are responsible for over 70% of contracts, and around 70% of contracts associated with a thread are linked to the top 30% of threads, suggesting a few influential members and threads play a prominent role. We denote top 5% influential members and threads contributing to the most contracts each *month* (both as maker and taker) as *key* members and *key* threads, respectively. Note that key members and key threads can be different for each single month.

Figure 6 shows the proportion of contracts made by key members and threads each month. While the proportion of key members over created contracts is consistently higher than over completed contracts, we see a opposite with key threads, which account for a higher proportion of completed contracts than created contracts. Except for created contracts of key threads, where we see a drop at the end of Set-up, the proportion of both key members and threads increases in Set-up, being more stable during Stable, and then drop at the end of this era. We observe a rapid increase of all types at the beginning of Covid-19, indicating the market became more centralised in response to the pandemic.

Network Centralisation. By looking at the underlying social graph made by the members, we find that the forum is also highly centralised around influential individuals in term of contractual connectivity. We consider two users n and m have a *raw* connection if they share at least one contract. An *inbound* connection is made from n to m if m accepts a contract from n while an *outbound* connection is made from n to m if n initialises a contract to m. Note that if a bidirectional contract (Exchange or Trade) is made from n to m or vice versa, both *inbound* and *outbound* connections will be taken into account for both n and m. The *raw*, *inbound*, and *outbound degrees* are the number of members that n has made at least a *raw*, *inbound*, or *outbound* connection with. These degrees reflect the connectivity representing the influence of members over the entire network. The higher the degree of n, the more influential n is in the marketplace.

Figure 7 illustrates the degree distributions (including *raw*, *inbound*, and *outbound*) of the market, over created and completed contracts (we omit plotting degrees that exceed 15). We observe

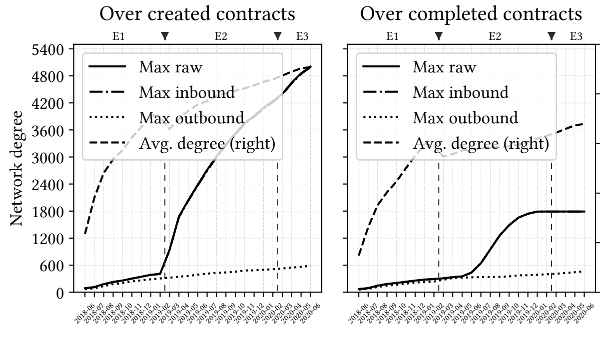

Figure 8: The growth of network degrees over time

for both created and completed contracts, for *raw* and *inbound* degrees, the network follows a power-law distribution, with most the nodes having few connections (mostly ranging from 1-15), however, there exists some nodes that are *extremely* linked (with *raw* degrees up to 5,004 for created and 1,790 for completed contracts, while *inbound* degree up to 4,992 for created and 1,789 for completed contracts). This distribution reflects a naturally grown scale-free network, which is different to randomly created ones, where the degrees of all nodes are typically distributed around the average.

For the *outbound* distribution, the highest *outbound* degrees are significantly smaller, at 587 created and 465 completed contracts. In other words, the most highly-connected nodes are formed mostly by *inbound* connections (accepting contracts) rather than making connection (initialising contracts) to others. For the most popular type – Sale, we observe many users initiating transactions, with a much smaller number of users accepting them. The *outbound* distribution is slightly different from *raw* and *inbound* as the number of nodes having zero connection is much smaller. Other than the zero point, outbound degrees also follow a power-law distribution.

Social Network Evolution. Figure 8 shows the average degree, max raw, inbound, and outbound degrees have increased over time. While the maximum outbound degree has gradually increased, the maximum raw, inbound, and average degrees have risen more dramatically. The max raw and max inbound are nearly identical, sharing the same pattern thoroughly (they overlap in the figure). While there was a gradual increase of all degree types during Set-up, there was a big uplift in Stable, when max raw and max inbound rocketed. Therefore, we believe the majority of raw degrees are made by inbound connections, suggesting the effect to the social network is mostly due to accepting transactions. No substantial change is observed during Covid-19, indicating the pandemic has not significantly affected the market in terms of contractual relationships. For both completed and created transactions the average degree grew gradually, suggesting a stable development of the contractual network. There was a slight drop in March 2019, which we believe is due to the sudden increase of the number of users, thus decreasing the average degree of the network. These new users also quickly contributed a large number of new contracts (see Figure 1).

4.3 Trading Activities

To classify trading activities, we first extract the obligation section in all *public* contracts, then apply normalisation techniques, such

Table 3: Number of completed *public* contracts (and unique users involved) in the top 15 trading activities

Category	Makers Side	Takers Side	Both Sides
currency exchange	5,533 (1,911)	5,281 (1,637)	9,516 (2,941)
payments	3,383 (1,415)	2,305 (1,122)	5,619 (2,198)
giftcard/coupon/reward	1,191 (367)	1,319 (484)	2,502 (744)
accounts/licenses	432 (230)	858 (435)	1,286 (598)
gaming-related	406 (194)	590 (265)	970 (409)
hackforums-related	544 (305)	486 (324)	911 (555)
multimedia	401 (50)	242 (70)	643 (95)
social network boost	194 (93)	372 (201)	564 (262)
tools/bots/software	129 (68)	384 (172)	508 (213)
tutorials/guides	153 (93)	339 (182)	485 (246)
hacking/programming	192 (101)	259 (190)	450 (263)
servers/rdp/domains	103 (51)	164 (101)	267 (136)
ewhoring	121 (25)	93 (48)	192 (64)
delivery/shipping	19 (4)	167 (33)	186 (36)
academic help	84 (30)	100 (43)	182 (62)
All Trading Activities	12,141 (3,211)	12,159 (3,313)	12,703 (5,135)

as removing stop-words, delimiters, digits, and unifying synonyms. We then use regular expressions to categorise trading activities into manually defined buckets. An *uncategorised* bucket is used for cases where the description is too short or ambiguously general to infer the category. Some of the categories are drawn from Motoyama et al. [16], while others have been added as they arose in the data, based on our domain-specific knowledge and from common goods observed from the text itself. We note that some contracts are placed in more than one category, for example, '*buying fortnite account*' would be categorised as both *gaming-related* and *account/license*.

Most makers initiate only a small number of contracts, with 49% making one transaction, 16% making two, and only 5% exceeding 20. Few makers account for the long tail, with just two users initiating over 700 contracts. Equally, most takers accept few contracts, with 46% responding to one, 16% to two, and 9% exceeding 20 contracts. The tail is longer for takers than makers, with two takers accepting more than 9,000 contracts. Therefore, although there is a small number of prolific users, most activity involves one-off transactions.

Products include automated bots, hacking tutorials, remote access tools (RATs), and eWhoring packs [18]. There is evidence the platform is being used as a cash-out market, with considerable activity trading giftcards, Bitcoin, and in-game currencies. Table 3 shows the top 15 commonly traded products and services, and number of unique users involved. We split the data by maker and taker to differentiate the parties involved. We note the total for all trading activities does not equal the sum of individual categories as some contracts fit into more than one category. For some transactions (e.g., exchanging currency), both sides are counted as one category, therefore the total is smaller than the sum of makers and takers.

A small number of categories account for the majority of trades with *currency exchange* and *payments* accounting for the largest proportion of contracts and unique users – *currency exchange* accounts for about 75% of all activities, over 70% higher than *payments*. *Giftcard* and *accounts/license* also account for a high proportion. In sixth place, *hackforums-related* refers to virtual Hack Forums products, such as buying *bytes* (a type of internal currency using within the forum), and *vouch copies* (the seller gives away the offered products for free, in exchange for vouches to gain more reputation). This

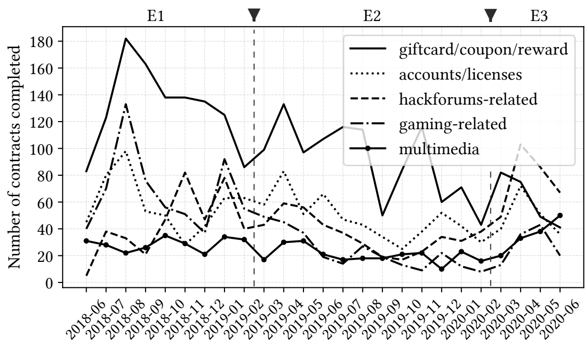

Figure 9: The evolution of top five products offered

leads further support to the idea that reputation plays a prominent role in the community. *Multimedia* (design, illustration and video editing) has the lowest proportion of unique members, indicating users tend to initiate and complete more repeated contracts for this category. The number of *delivery/shipping* takers are seven times higher than makers, demonstrating an imbalance in demand. We see *academic help*, which includes assistance with homework, essays and dissertations, accounts for a small proportion of transactions.

The Evolution. Figure 9 shows the evolution of the top five *products* over the three eras. We exclude *currency exchange* and *payments* here and separately examine these in §4.4. Overall, *giftcard* is the most popular traded product in all three eras. *Gaming-related* peaks during Set-up, but drops to the lowest position by the end of Stable. *Hackforums-related* also grows in Set-up, but slips back to fourth position by the end of this era. In Stable, *accounts/licenses* consistently take the second-highest position. Despite the increased number of contracts at the beginning of Stable (see Figure 1), the number of contracts in the top five categories does not grow rapidly – some even decrease over time. This is likely due to the decline in the number of *public* contracts (see Figure 2), with users moving to *private* contracts where the details are hidden. We see a rapid stimulus of all activities during Covid-19. While there is a consistent increase of *multimedia*, the others perform a rapid but short-lived peak. At the end of the era, *hackforums-related* comes to take the highest position, despite placing last at the beginning of Set-up, suggesting a high demand for reputation, while *multimedia* takes the second position and *giftcard* drops to its lowest ever position.

4.4 Payment Methods

As *currency exchange* and *payments* account for most trading activities, we take a closer look at the types of currencies and payment methods used. We first take all the contracts classified into *currency exchange*, *payments*, and *giftcard*, then apply another regular expression set to categorise the payments used. As shown in Table 4, Bitcoin and PayPal are the most popular payment methods, accounting for 75% and 38% of completed contracts, respectively. Amazon Giftcards are ranked third. The most wanted fiat is USD while JPY, GBP, EUR, and CAD account for a tiny proportion of transactions (not shown in the table). Other cryptocurrencies, including Ethereum, Bitcoin Cash, Litecoin, and Monero, account for trivial proportions, indicating that despite its limitations, Bitcoin is still a popular cryptocurrency on the underground marketplace.

Table 4: Number of completed *public* contracts (and unique users involved) in the top ten payment methods

Payment Methods	Makers Side	Takers Side	Both Sides
Bitcoin	4,456 (1,646)	4,486 (1,416)	8,787 (2,559)
PayPal	2,561 (1,202)	1,926 (976)	4,465 (1,908)
Amazon Giftcards	986 (287)	771 (279)	1,754 (493)
Cashapp	559 (204)	209 (126)	767 (304)
USD	196 (131)	350 (140)	543 (263)
Ethereum	230 (147)	138 (103)	362 (224)
Venmo	134 (65)	74 (48)	206 (104)
V-bucks	95 (13)	65 (15)	159 (19)
Zelle	77 (43)	44 (22)	121 (62)
Bitcoin Cash	20 (18)	24 (23)	44 (38)
All Methods	9,358 (2,802)	8,058 (2,450)	11,793 (4,276)

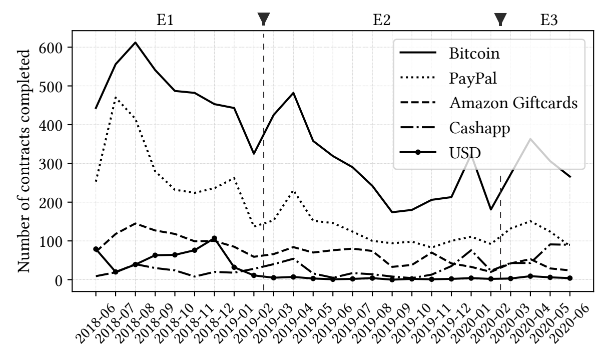

Figure 10: The evolution of top five preferred payment methods used in completed *public* contracts

The highest rate of repeat transactions is of V-Bucks (Fortnite's in-game currency), with 8.37 transactions per trader.

The Evolution. Figure 10 shows the evolution of top payment methods in *completed* contracts. Bitcoin and PayPal dominate in all three eras, confirming findings in prior research [19, 22]. Amazon Giftcards retain third place most of the time. After peaking in the early stage of Set-up, the number of contracts for the top three payment methods gradually declines. USD surpasses Cashapp in Set-up, then drops to fifth position, where it remains. There is a gradual downtrend in Stable, despite the increased number of contracts and members at the beginning of this era (see Figure 1). Again, we believe this is because many users chose *private* transactions when contracts became mandatory (Figure 2). In Covid-19, there is a short-lived rise in contracts for all payment methods, particularly Bitcoin and PayPal. At the end of this era, Cashapp outpaces PayPal and Amazon Giftcards to second place, its highest ever ranking.

4.5 Trading Values

We estimate the trading values for *completed* contracts, ignoring Vouch Copy (as they are proofs of reputation rather than an economic trades). We apply regular expressions to extract the trading values and currency denominations quoted in the maker/taker obligation sections. Note the extracted values are from the *contractual arrangements*, rather than actual transactions. The trading volume

Table 5: Top 10 trading activities and payment methods by contracts values

Trading Activities	Value (Makers)	Value (Takers)	In Total	**Payment Methods**	Value (Makers)	Value (Takers)	In Total
currency exchange	$522,125	$449,103	$971,228	Bitcoin	$419,395	$389,888	$809,283
payments	$274,836	$199,723	$474,559	PayPal	$177,512	$156,913	$334,425
giftcard/coupon/reward	$49,156	$69,089	$118,245	Amazon Giftcards	$45,511	$59,735	$105,246
hacking/programming	$18,291	$11,454	$29,745	Cashapp	$68,597	$15,029	$83,626
accounts/licenses	$13,099	$9,064	$22,163	USD	$15,014	$36,771	$51,785
social network boost	$7,210	$9,960	$17,170	Venmo	$14,777	$11,986	$26,763
tutorials/guides	$3,145	$13,051	$16,196	Zelle	$10,019	$12,121	$22,140
marketing	$2,750	$4,872	$7,622	Ethereum	$14,312	$7,461	$21,773
contest/award	$1,015	$3,070	$4,085	Apple/Google Pay	$3,157	$534	$3,691
tools/bots/software	$2,245	$1,610	$3,855	Bitcoin Cash	$2,031	$519	$2,550

is then estimated by counting all extracted values naively, as there is no way to confirm if transactions actually went through as described. We assume any goods without a value and denomination specified, such as 'dissertation', have an equal value with the opposite side. If the values of both sides cannot be estimated, the contract value is ignored. If a transaction results in double counting, where there are values observed in both the maker and taker sides (e.g., currency exchange), we take the average as the final value. If no specific denomination is declared, or it cannot be inferred from the text, we consider it to be USD, as this is the currency most transactions are traded in (followed by GBP, CAD, EUR, AUD, and INR). We then convert all values to USD using the conversion rates at the time the transactions were made.

To estimate the trading values more precisely, we manually check the 163 high-value (exceeding 1,000 USD) transactions, which is found mostly related to Bitcoin and PayPal (or Cashapp) exchanges. We then verify these by manually reading the obligations, terms, payments information, ratings, and advertising threads (if any) to identify actual values. For contracts providing a Bitcoin address and/or transaction hash, we additionally check recorded transactions on the blockchain at the completion time. If we can not determine values for both sides of the contract, it is excluded. We found in many cases values exceeding $10,000 are likely due to typing errors. Of the 163 high-value trades, 82 (50%) are confirmed, 11 (7%) could not be confirmed, and 70 (43%) have a different (usually lower) trading value. We also see indications of private negotiations, for example, one service is advertised at $1,000 but the actual transaction (verified on the blockchain) is $200, In some cases the transaction value on the blockchain is higher, for example, a contract stating $1,250 Bitcoin but a transaction value of $1,800. We also observe Bitcoin is often traded with a higher value than itself, for example, $1,000 BTC for $1,080 PayPal, indicating Bitcoin is probably in high demand compared to other cashout methods. We manually update the contract details based on the new values observed.

The total value of *public* transactions across contracts with non-zero values is estimated to be $978,800 (average $85, maximum $9,861). Within each type, this corresponds to $461,484 in Exchange (av. $104, max. $9,000), $304,783 in Sale (av. $71, max. $6,723), $205,247 in Purchase (av. $78, max. $9,861) and $7,286 in Trade (av. $58, max. $400). Table 5 show the values ranked by top 10 trading activities and payment methods. Note that these results are *naively* calculated by summing the value of each activity. Thus, their totals are higher than the total value of transactions estimated above, as some transactions are classified into multiple categories.

With regards to trading activity, *currency exchange* accounts for the highest value ($971,228), nearly double the second highest, *payment*, followed by *giftcard*. They are also the top three trading activities by number of contracts as shown in Table 3. By payment methods, Bitcoin accounts for the highest amount ($809,283), nearly 2.5 times higher than PayPal ($334,425). A small number of members are involved in a large proportion of the total value traded, with the top 10% users party to over 70% of the total value. This again demonstrates a high centrality of key members in the marketplace. On average, a user who is party to at least one contract makes around $185 of trading value during the entire period.

The actual trading values are likely to be much larger, as the proportion of completed private contracts is over five times higher than public ones which for the past year have accounted for only 15.7% of transactions. To estimate the value for both private and public contracts, we assume private contracts are *at least* as valuable on average as public ones. We note more valuable transactions accord a higher degree of risk of incrimination, and therefore may be more likely to be private. One way to see the representation of public transactions is looking at the *disputed* contracts, in which the contract details became *public* from *private*. Among disputed contracts, while most users are only involved in one dispute, we observe one user has a relatively high record with 21 disputes. Otherwise, we do not see any abnormal and questionable behaviour in term of trading goods and services, with most disputed transactions exchanging Bitcoin, and some relating to eWhoring. We thus extrapolate by *each contract type* to gain a lower bound total estimated value of $6,170,943 for both public and private contracts.

The Evolution. Figure 11 shows the evolution of monthly value by contract types, top five payment methods, and top five product categories traded over the three eras. Exchange generally accounts for the highest value, followed by Sale and Purchase. Trade consistently accounts for the smallest value. The value of Exchange during Set-up declines from its peak in July 2018 to the lowest position in this era in February 2019. After peaking during Set-up, Sale and Purchase also gradually decrease towards the end of Stable. In Covid-19, we observe a short-lived increase in Sale, which for a while outpaces Exchange to become the highest value in March and April 2020, however Exchange quickly resumes first place afterwards. Purchase and Trade values, on the other hand,

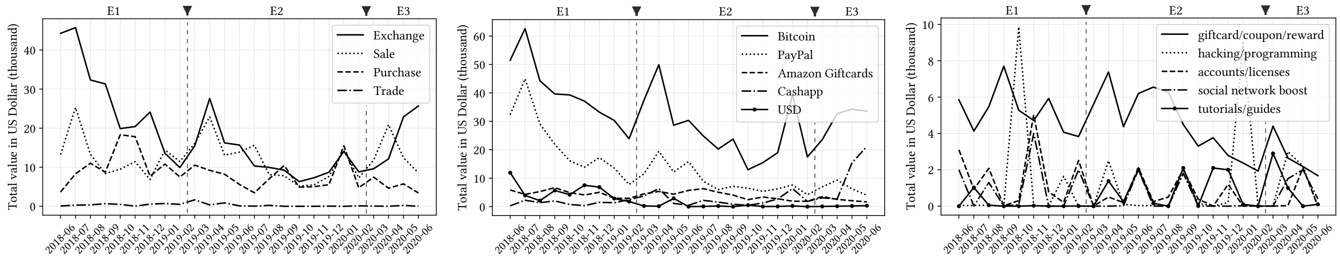

Figure 11: The evolution of monthly value made by contract types, top five payment methods and top five products

did not increase much, showing a very high demand of Exchange and Sale compared to others.

Bitcoin and PayPal dominate the other payment methods in all three eras. However, the difference in value is considerably large. During Set-up, the value made by Bitcoin and PayPal dropped quickly, from its peak on July 2018, to its lowest point in this era on February 2019, while others stayed at around $9,000-10,000 per month. During Stable, there is no significant change in the order of these methods, except for Cashapp in April 2019 and January 2020, when it outpaces Amazon Giftcards and USD, becomes the third-highest method, then slips back quickly. The most fluctuating method during this era is Bitcoin, with two peaks in April 2019 and January 2020. During Covid-19, the monthly values increase across most payment methods, with Bitcoin climbing around 90% higher compared to the end of Stable. In this era, we see Cashapp outpacing PayPal, Amazon Giftcards and USD to become the second highest wanted payment method. The largest observed difference between Bitcoin and PayPal is at the end of this era (June 2020), when total Bitcoin value is eight times higher than PayPal.

There is a more complex fluctuation with regards to products traded. Overall, *giftcard* is consistently the top category. In Set-up, we observe an increase of *hacking/programming* in October 2018, reaching nearly $10,000. This was due to some high value contracts, which we manually checked and concluded were actual trades. The peaks for *accounts/licenses* and *social network boost* also happen in November 2018. During Stable, *giftcard* gradually declines, reaching the lowest value at the end of this era. The others change frequently, but not significantly, ranging mostly under $2,000, except a sudden lift of *hacking/programming* on January 2020. We manually checked this increase, again confirming it is likely to be correct. We observe a raise then decline in all top five products during Covid-19. While *giftcard* and *tutorials/guides* are immediately affected in the first month of this era, *hacking/programming* and *social network boost* stay mostly unchanged in the first couple of months before peaking (then dropping) in the month after.

5 LONGITUDINAL ANALYSIS

In this section, we use statistical modelling approaches to examine the longitudinal evolution of the marketplace through the three eras then address the 'cold start' problem.

5.1 Latent Class Transition

To generate insights into the kinds of users and patterns of behaviour in the market, we use Latent Transition Modelling (LTM) [14] to identify latent 'classes' within the data. LTM involves a longitudinal application of Latent Class Analysis, a statistical modelling technique which uses clustering to find latent groups in data which share similar characteristics, and to assign group membership to the items in our dataset. In this case, we classify users at a point in time, based on the number of transactions they make of different kinds. This is crucial to understanding how co-operation evolves over time in this marketplace and allows us to assess the representativeness of our data and the structure of the market – for example, whether the activity we observe represents a large number of small-time users or is dominated by a few key players.

By creating a Latent Transition Model, we can additionally understand how users move between classes over time and how they change across the lifetime of the market. The model treats each month's activity for each user as a separate case. Using a Poisson curve (due to non-overdispersed count data), the most accurate and parsimonious (per AIC and BIC) is a 12-class model. Effectively, this distils the complexity of market activity down to 12 types of users (summarised in Table 6) and assesses their contribution to different types of activity.

We apply the model to the three eras, focusing on the contribution of different classes to marketplace activity over time (rather than the total *number* of individuals in each class). Thus, we report on the *total number of transactions* over time made by users exhibiting different classes of behaviour. This allows us to explore, for example, the proportion of Exchange made by a small number of 'big fish' versus those made by large numbers of bit players at a given period in time. To establish the links between makers and takers, we also report on 'flows' in each of the eras (summarised in Table 8 in the Appendix), providing details of the three maker-taker pairs accounting for the highest percentage of each contract type in a given era. We ignore Trade and Vouch Copy as they account for small proportions of contracts. We see low levels of disputed transactions (around 1%) for most of this period, but these peak to 2-3% for the last six months of Set-up.

The Set-up Era. The class distribution of those making and accepting Exchange is roughly similar, with most classes who make contracts also accepting contracts made by others. At first, the exchange market is dominated by a large number of users who participate in only one or two transactions (Figure 12). Around two-thirds of Exchange involve these users, with the remaining third involving 'power-users'. After the first six months, the growth in Exchange is largely driven by small numbers of power-users, who by the end account for the majority of transactions. In this era,

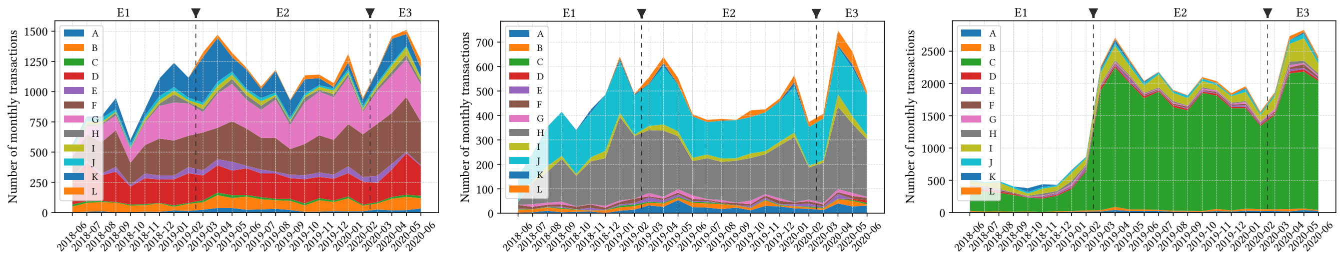

Figure 12: Number of Exchange/Purchase/Sale transactions which were *made* by different classes over time

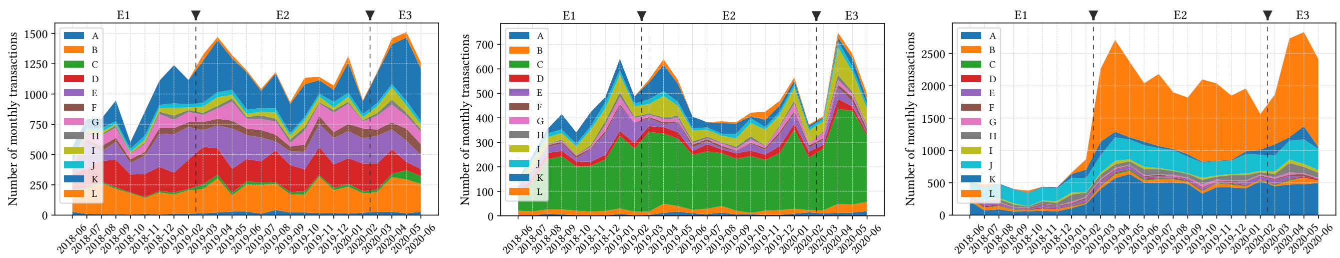

Figure 13: Number of Exchange/Purchase/Sale transactions which were *accepted* by different classes over time

Table 6: The average numbers of monthly transactions made over latent classes

	Make					Accept					
	Exchange	Purchase	Sale	Trade	Vouch Copy	Exchange	Purchase	Sale	Trade	Vouch Copy	Behaviour type
A	0.5	0.6	0.5	0.1	0.0	0.5	0.2	10.1	0.2	0.0	Mid-level Sale taker
B	2.3	0.4	0.6	0.1	0.0	6.5	0.6	1.1	0.1	0.0	Exchanger & Sale taker
C	0.0	0.0	1.1	0.0	0.0	0.0	0.2	0.0	0.0	0.0	Single Sale maker
D	0.9	0.0	0.1	0.0	0.0	0.9	0.1	0.0	0.0	0.0	Single Exchanger
E	4.3	0.7	2.0	0.2	0.0	22.3	4.2	3.8	0.4	0.0	Exchanger power-user
F	7.3	0.2	0.4	0.0	0.0	1.3	0.2	0.3	0.0	0.0	Mid-level Exchanger
G	21.2	0.6	1.3	0.1	0.0	8.1	1.1	1.3	0.1	0.0	Exchanger power-user
H	1.3	10.0	0.9	0.2	0.0	1.0	0.4	3.2	0.1	0.0	Mid-level Purchase maker
I	1.1	0.7	5.2	0.2	0.0	1.6	2.0	1.0	0.1	0.0	Mid-level Sale maker
J	0.1	0.7	0.1	0.0	0.0	0.1	0.1	1.1	0.0	0.0	Single Sale taker
K	31.2	0.9	3.3	0.3	0.0	54.9	9.2	12.8	1.0	0.0	Exchanger power-user
L	1.3	1.1	1.2	0.2	0.1	1.5	0.6	54.9	0.2	0.0	Sale taker power-user

power-users and single exchangers are not well connected, with most flow volumes trading within their own class types.

For Purchase, we observe a different story. Overall transaction volume ramps up steadily across Set-up. In contrast to Exchange, there is a clear division between 'maker' and 'taker' classes, with little overlap. The majority of Purchase are *made* by two classes: H (a medium-class making on average 10 Purchase a month, with some participation in Exchange); and J (a low-volume class making around 1 Purchase and accepts 1 Sale every month). Conversely, accepting Purchase is dominated by low-volume users from class C (who accept a single Purchase per month on average) (Figure 13). This changes over the last six months of Set-up, with a rise in prominence of class E in accepting transactions (one of our exchanger power-user classes).

We again observe very different class participation in making and accepting Sale. The majority of transactions are made by class C (small-scale users who make a single Sale per month on average and nothing else). In Set-up, takers of Sale are dominated by class J (small-scale users who accept a single sale transaction). This arrangement, of large numbers of small-scale users selling to one another on a one-to-one basis and in a relatively focused manner, continues for most of this era.

The Stable Era. The transition between Set-up and Stable sees a rush of small-scale users making small numbers of Sale. An initial peak after contracts became compulsory is followed by a slow downward trend. We see a small spike in Purchase and Exchange around Christmas/New Year 2019.

Exchange remains relatively stable, with the class mix for making and accepting continuing to reflect the end of Set-up. There is a rapid transition at the beginning of Stable, with Sale roughly quadrupling in volume over three months for class C users (small-scale users making a single sale). On the taker side, we see a starker change in the users accepting Sale. While small-scale exchangers accept about the same number as they did during Set-up, two new classes emerge: L (power-users accepting on average 54 Sale per month); and A (a medium-user class accepting 10 Sale per month).

The distribution of Purchase changes little during the Stable era. The main change relates to takers, with a short-term initial growth in classes I (a medium-scale class with an interest in a variety of transaction) and K (an Exchange power-user class that also accepts medium numbers of Sale and Purchase) which drops off after four months. For Sale, we observe a sustained increase in users making single transactions being accepted by a large increase in a new class, L (Sale power-users) and an existing one, A (Sale medium-level users), with most of the market being split between these two. We see disputes drop at the start of this era to around half or third their previous proportion of contracts (around 1%).

The Covid-19 Era. This era heralds changes to the market, with a large increase in all three contract types (Exchange, Purchase, and Sale). There is a large influx in class C users, who make a single Sale. We observe increases in contracts across classes B, D, G, and L, all of which are small-volume exchange classes apart from G. Increases in accepted Exchange are concentrated in classes B and K (exchanger power-users). For Purchase, increases in contracts made are split between small-scale (A class) users and mid-level users (H), while for Purchase accepted, increases are across classes C, B, D, and E (all small-scale users). Increases in Sale are focused in two classes: small-time users (class C) making Sale and sale power-users (class L) accepting them.

5.2 Cold Start Problem

New members – *cold starters* – face the challenge of getting started on the market, establishing reputation and building up a customer base. We use clustering and qualitative analysis to investigate how users in Stable overcome the 'cold start' problem. We then use Zero-inflated Poisson regression to explore the role of trust and reputation for cold starters across all three eras.

Cold Start Variables. For modelling, our choice of predictor variables is informed by the literature on trust in underground markets, and include users' positive and negative ratings, number of disputed transactions, and length of participation since first active post [11, 16, 27]. Included as control variables are the number of posts in the marketplace and the number of contracts initiated and accepted, to control for users' level of activity and social ties [7, 11]. For clustering, we use the predictor and control variables.

Clustering. We use k-means clustering and the *cold start variables* to examine groups within a subset of members who accepted their first contract in Stable, our second era of analysis. We limit this analysis to Stable, as during Set-up many actors had a presence in the marketplace before the contract system began. We find two clusters are the best fit for our dataset. The first cluster contains the majority (97.7%) of members. These users have a median of one accepted contract and seven posts. The second cluster is significantly smaller, containing 2.3% of members, with a median of 49 accepted contracts and 279 posts. Thus, this cluster is characterised by a greater amount of market activity.

Despite standardising our variables (zero mean and variance of 1) to give them equal weight, our dataset is skewed, which has implications for the clustering results. As k-means clustering relies on distances, it is expected one cluster will be significantly larger, representing the *general low-volume activity* of most members. However, it is the *outliers* we are interested in, as these capture both the users who successfully overcame the 'cold start' problem, and users who have significantly different activity to most members, such as a higher number of disputes. To explore these outliers, we again cluster this group of 122 members, finding eight clusters (see Table 7). We use cluster G, with the highest proportion of members, as the *baseline* group for comparison.

Table 7: Each cluster in the outlier group, with size and median variables (+ for positive feedback, − for negative feedback, MPosts for posts in the marketplace, Maker/Taker for contracts made/accepted)

	Size	Disputes	Posts	+	−	MPosts	Maker	Taker
A	12	0.0	1930.5	10.5	0.0	137.0	17.0	17.5
B	29	0.0	86.0	54.0	0.0	55.0	29.0	157.0
C	2	14.5	205.5	74.5	7.5	139.5	36.0	257.0
D	8	0.0	192.5	35.5	0.0	117.0	196.0	50.0
E	2	0.0	219.5	240.0	1.0	151.5	109.0	485.5
F	5	1.0	944.0	135.0	0.0	698.0	99.0	203.0
G	43	0.0	491.0	19.0	0.0	204.0	47.0	38.0
H	21	4.0	211.0	10.0	1.0	131.0	15.0	27.0

Clusters B and G have the highest proportion of members, with B having a lower marketplace post count, and higher number of accepted contracts and positive ratings. Cluster A has the highest number of post counts, but a lower number of accepted contracts than the baseline. Cluster F has more marketplace posts, while cluster C has a significantly greater number of disputes and negative ratings, and cluster H has the second most disputes and negative ratings, with more members than cluster C. Cluster D has the most initiated contracts, indicating these members are making as well as receiving contracts. Cluster E has the most positive ratings and accepted contracts.

The median lifespan of activity of new users on the contract system in Stable is less than one day, while for the group of outliers, it is 250 days. Of all cold starters in Stable, 13.0% of members, and 54.1% of members in the group of outliers, continue accepting contracts into Covid-19. Looking at *reputation voting*, the median reputation score for cold starters in Stable is 33, while it is 157 for the outlier group. Members starting in Set-up have a median reputation score of 96, which is greater than starters in Stable. This is due to more members in Set-up having an existing presence in the marketplace, before the contract system started.

Types of Products and Services. As Wegberg et al. [25] found product characteristics influence sales volumes, we investigate the types of products or services being marketed. We qualitatively analyse threads associated with public transactions, which serve as advertisements for a given product or service, to reveal the extent product type drives trade. For this analysis, we only consider the *completed* contracts of high volume users in Stable. We find that in the context of the 'cold start' problem, the transaction type plays an important role. The majority of these members build their reputation by participating in Exchange, where a product is exchanged for another item. A proportion of these users only offer items on an exchange basis, while the remainder are involved in Sale and other types of contracts. Most Exchange are for currency exchange, such as PayPal to Bitcoin (and vice versa), Ethereum to Bitcoin, PayPal to Apple Pay, Bitcoin to Cashapp, and Bitcoin for Giftcards. A small proportion of users do not participate in Exchange, instead establishing themselves by offering products and services. These include eWhoring packs and tutorials, the 'YouTube method' (usually tutoring in basic passive income schemes such as dropshipping), hosting, botnets, and software upgrades/licenses.

Trust and Reputation. To understand the role of trust and reputation in the cold-start problem, we model completed contracts using Zero-Inflated Poisson (ZIP) models and the *cold start variables* (including controls). The ZIP models provide statistical estimates on two processes: the expected number of completed contracts (the count model) and the odds of having zero completed contracts (the zero-inflation model) for users in the contract system. We run the full-sample models for all eras and the sub-sample models (first-time users and existing users) for Stable and Covid-19. Given their skewed distributions, all variables, with the exception of length and number of completed contracts, are transformed using the square-root function. We note the variables for all count models are measured for each era. Results from Vuong tests for all models suggest the ZIP models are better-fitted for the data.

During Set-up, there are 6,278 users of the new contract system. In general, the more active users are with the contract system during Set-up, the more completed contracts they had. For example, an increase of one post in the marketplace increases the expected number of completed contracts by 1.04. The *zero-inflation model* shows that users' negative ratings and length of activities lower the odds of having zero completed contracts by 0.578 and 0.991 respectively. Disputes also lower the odds, but are not statistically significant. The results suggest that despite negative feedback from other users, first adopters of the contract systems who were active were successful in completing transactions during Set-up.

In Stable, with 16,123 first-time contract users and 3,534 existing users, the *count model* shows comparable results to Set-up, suggesting that active users in the contract system continue to have success with transaction completion. Being first-time users did decrease the number of expected contracts but they were not penalised in terms having any completed contracts, as indicated in the *zero-inflation model* in Table 9. The sub-group results show that first-time users who received more negative ratings and had disputed contracts had lower number of completed contracts and increased odds of zero completed contracts, respectively. Existing users are not penalised with the same cautiousness, as having negative ratings and disputes respectively increased the number of completed contracts and lower the odds of having zero completed contracts. The results suggest first-time users of the contract system were treated with suspicion during Stable.

During Covid-19, there are 2,569 first-time users and 5,275 existing users of the contract system. Both the *count and zero-inflation models* with all users yields comparable results to the Stable era. First-time users continue to have fewer completed contracts but had lowered odds of zero completed contracts. As a sub-group, first-time users with more positive rating and posts in marketplace had lower expected number of completed contracts. This differs from Stable where trustworthiness was assessed based on negative feedback and disputes. These effects are, however, absent for existing users. The effect of disputes on lowering odds of zero-completed contracts remain for existing users during Covid-19. The findings suggest that first-time users are held to different standards of trustworthiness then existing users.

6 DISCUSSION

The Hack Forums marketplace provides a range of trust capabilities to facilitate trade between pseudonymous parties. By having a semi-public record of all transactions, the marketplace affords users a trust infrastructure, which allows new users with no established reputation build up trust through making initial small-scale exchanges which are publicly recorded. As with other contemporary online markets, *measurement* has become a core part of trust. Tracked transactions and semi-visible histories are signals of reputability that go beyond a single reputation number or patchy list of feedback. Given the majority of transactions are private, with only minimal details provided, the overall outcome of this new marketplace is to further centralise control (under the guise of a *public* trust mechanism) to the forum administrators. Having administrators act as third-party arbitrators is similar to the mediation model in many legitimate marketplaces, such as eBay or Amazon, where disputes are mediated by the platform. This increases trust in the intermediaries, particularly buyers' trust toward the marketplace itself. In addition to studying the trust functions facilitated by this marketplace, we track its evolution over three main eras.

The Set-up Era. The market forms in the first era, with users gradually shifting to the new platform. Initially Exchange contracts are split between large numbers of small-scale users (who make single currency exchanges) and power-users. We see an increase in Exchange, largely driven by power-users who then dominate this contract type. Purchase is dominated by 'small-fry', with the growth in transactions almost entirely driven by single transaction users, with takers split between small-fry and a small number of power-users. Sale is dominated by small-scale users in this era. Although this period begins with an even mix of public and private transactions, it shifts in favour of private transactions towards its end, when contracts become compulsory. We do not see the growth in this era of a 'concentrated' market, with small-scale users dealing with other small-scale users, and power-users with other power-users. Thus, as the market slowly grows, it is not turning into a 'business to customer' market. This has important implications for trust mechanisms across Set-up, as the trust function of recording transactions is likely to play only a small role where individuals make or accept only a single transaction. Thus, in Set-up, we conclude that the market largely facilitates the growth of relationships between power-users, rather than the establishment of trusted traders used by large numbers of small-scale users.

Although Tuckman's theory of group development relates to group interactions with a common goal, rather than competitive entrepreneurial markets, its general contentions with respect to group formation are useful, given sustained facilitation of trust and exchange is a *group endeavour* in which all users (apart from scammers) are involved. In particular, when conceived more broadly as pertaining to social interaction, Tuckman's theory draws out useful longitudinal aspects of *conflict* and *consensus.* In the forming era, with groups coming together, and the subsequent 'storming' phase, aspects of intra-group conflict emerge that need to be resolved. Thus, we argue that this gradual and dispersed Set-up era reflects individuals testing the system and establishing their orientations with respect to one another (as we see from the super-user to super-user transaction patterns). We also see evidence of the 'storming' phase at the end of this era, with a spike in disputes before the transition to Stable.

The Stable Era. There is a considerable shift in the composition and scale of the market when contracts become compulsory. While far more transactions are being made, the majority (around 88%) are now private, meaning other users can only see limited feedback and the transaction type. The market sustains a diverse range of behaviours and products over this period. We observe the growth of 'business to customer' patterns of trade, with individual power-users beginning to cultivate large numbers of small-scale customers, rather than trading with one another.

The influx of customers appears to accelerate competition (and hence, conflict), accelerating progression to the *norming* phase, our Stable era. Power-users who established themselves over the Set-up period capitalise on the reputation and trust they have built up, and alongside newer would-be power-users, are well-positioned to capitalise on the influx of small-scale custom. Hence, the trust relationships facilitated by the market infrastructure shift between these two phases - from a forming, orienting function to one which more closely resembles trust relationships within a traditional market, with clear producers and consumers.

The Covid-19 Era. The most timely aspect of this paper relates to the initial effect of the Covid-19 pandemic on the market for cybercrime products and services. Across all our measures, we conclude that the effects of Covid-19 are a *stimulus* rather than a *transformation*. The same kinds of transactions, users, and behaviours dominate as during Stable, however volumes increase for all product categories. In terms of users, we see an uplift in numbers across most categories, but particularly in users making small numbers of Sale. In the short term, this returned the market to the state it was in at its previous peak (from which it had been steadily declining). We observe increases in currency exchanges, and most other product categories. Much of the increase in transactions is due to increased *flows* of merchandise and currency between small-scale actors and power-users, suggesting that the crisis has concentrated the market around these power-users. Thus, it is largely existing power-users who benefit from the opportunities provided by the Covid-19 crisis.

We believe that the most convincing explanation for this is *situational*. Rather than serious new forms of crime arising, the uplift may result from changes in the everyday lives of the people who use this market. From consulting forum posts made on Hack Forums during this period, the picture is one of mass boredom and economic change. Younger users are confined to their homes with no school and an abundance of time, while older users have either been laid off work and are time-rich and money-poor (and hence desperate to sell and exchange) or are spending their recently-received stimulus cheques. Thus, coronavirus is 'turning up the dial' on the factors already feeding into marketplace participation. We also note that few products require shipping physical goods, so the trust mechanisms are not being strained by the lockdown in the way that they might be on, for example, drug-related cryptomarkets.

7 CONCLUSION

This is the first academic study, of which we are aware, on the evolution of an underground marketplace (including its shift through the Covid-19 pandemic) using the contractual transactions made and completed by the forum's members. We have used quantitative analysis and statistical modelling approaches to outline the economic shape of the forum's market, including the sorts of goods and services being exchanged, the money being made, preferred payment methods used, changes in the market over time, and how users overcome the 'cold start' problem when joining the market without established trust and reputation. We conclude that the contract system constitutes a useful trust infrastructure for participants.

The broader relevance of this paper to the security community is threefold. First, our analysis of this novel dataset suggests that centrally-held, mostly-public records of transactions provide a form of trust and reputation infrastructure which appears to particularly benefit the concentration of the market over time around a core of power-users. This has implications for intervention, particularly for approaches that involve confusing the 'trust signals' which make up this public record. We suggest spurious negative reviews and other forms of Sybil attack are best targeted in the early days of market formation, before this concentration effect takes root and while trade is largely between parties of similar size. Second, our analysis represents a novel application of clustering methods to a database of illicit transactions, demonstrating the usefulness of traditional statistical modelling techniques such as LTA for tagging and labelling large administrative datasets, facilitating data reduction and 'data science' analysis. Third, our analysis suggests the pandemic has had an effect on the market, but mainly due to an influx of small-scale customers which has largely benefited existing power players, rather than enabling small-scale sellers to make a jump to the 'big leagues'.

Our dataset has some limitations. First, while contractual details can be observed, we generally have no way to verify if transactions actually go through with exact values as described. Where traders specify the Bitcoin/Ethereum addresses or the transaction hashes, the actual values on the blockchain can be confirmed, otherwise, the dataset lacks ground truth verification. Moreover, even when the transaction hash is provided, we have no way to verify its integrity, as the dishonest parties could still find an appropriate transaction on blockchain then put it into the contract details to gain reputation. Second, the proportion of private contracts in our dataset is considerable (around 88%), which means some details are hidden from us. Our dataset is made available for academic research through a data sharing agreement. We hope our analyses presented in this paper, as well as the dataset we provide, enable further studies into better understandings of how underground economies have been established, operated, and evolved.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers and our shepherd Zakir Durumeric for their valuable and insightful comments. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC)[grant numbers EP/M020320/1 and EP/V026178/1]. The conclusions and findings presented in this work are those of the authors and do not necessarily reflect the opinions of the funders.

REFERENCES

[1] Sadia Afroz, Vaibhav Garg, Damon McCoy, and Rachel Greenstadt. 2013. Honor among thieves: A common's analysis of cybercrime economies. In *2013 APWG eCrime Researchers Summit*. IEEE, New York, 1–11.

[2] Luca Allodi, Marco Corradin, and Fabio Massacci. 2016. Then and now: On the maturity of the cybercrime markets the lesson that black-hat marketeers learned. *IEEE Transactions on Emerging Topics in Computing* 4, 1 (2016), 35–46.

[3] Rasika Bhalerao, Maxwell Aliapoulios, Ilia Shumailov, Sadia Afroz, and Damon McCoy. 2019. Mapping the Underground: Supervised Discovery of Cybercrime Supply Chains. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, New York, 1–16.

[4] British Society of Criminology. 2015. Statement of ethics. http://www.britsoccrim.org/ethics/

[5] Andrew Caines, Sergio Pastrana, Alice Hutchings, and Paula J Buttery. 2018. Automatically identifying the function and intent of posts in underground forums. *Crime Science* 7, 1 (2018), 14.

[6] Nicolas Christin. 2013. Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace. In *Proceedings of the 22nd International Conference on World Wide Web*. ACM, New York, 213–224.

[7] David Décary-Hétu and Anna Leppänen. 2016. Criminals and signals: An assessment of criminal performance in the carding underworld. *Security Journal* 29, 3 (2016), 442–460.

[8] David Gefen, Elena Karahanna, and Detmar W Straub. 2003. Trust and TAM in online shopping: An integrated model. *MIS Quarterly* 27, 1 (2003), 51–90.

[9] Thomas J. Holt. 2013. Examining the forces shaping cybercrime markets online. *Social Science Computer Review* 31, 2 (2013), 165–177.

[10] Thomas J. Holt and Eric Lampke. 2010. Exploring stolen data markets online: products and market forces. *Criminal Justice Studies* 23, 1 (2010), 33–50.

[11] Thomas J Holt, Olga Smirnova, and Alice Hutchings. 2016. Examining signals of trust in criminal markets online. *Journal of Cybersecurity* 2, 2 (2016), 137–145.

[12] Thorsten Holz, Markus Engelberth, and Felix Freiling. 2009. Learning More about the Underground Economy: A Case-Study of Keyloggers and Dropzones. In *Computer Security – ESORICS 2009*. Springer, Berlin, 1–18.

[13] Xuan Nhat Lam, Thuc Vu, Trong Duc Le, and Anh Duc Duong. 2008. Addressing Cold-Start Problem in Recommendation Systems. In *Proc. 2nd Int. Conf. on Ubiquitous Information Management and Communication*. ACM, New York, 208–211.

[14] Paul Felix Lazarsfeld and Neil W Henry. 1968. *Latent Structure Analysis*. Houghton Mifflin Co., Boston.

[15] Blerina Lika, Kostas Kolomvatsos, and Stathes Hadjiefthymiades. 2014. Facing the Cold Start Problem in Recommender Systems. *Expert Syst. Appl.* 41, 4 (2014), 2065–2073.

[16] Marti Motoyama, Damon McCoy, Kirill Levchenko, Stefan Savage, and Geoffrey M. Voelker. 2011. An Analysis of Underground Forums. In *Proceedings of the Internet Measurement Conference*. ACM, New York, 71–80.

[17] Meltem Odabaş, Thomas J Holt, and Ronald L Breiger. 2017. Markets as governance environments for organizations at the edge of illegality: insights from social network analysis. *American Behavioral Scientist* 61, 11 (2017), 1267–1288.

[18] Sergio Pastrana, Alice Hutchings, Daniel Thomas, and Juan Tapiador. 2019. Measuring eWhoring. In *Proceedings of the Internet Measurement Conference*. ACM, New York, 463–477.

[19] Sergio Pastrana, Daniel R. Thomas, Alice Hutchings, and Richard Clayton. 2018. CrimeBB: Enabling Cybercrime Research on Underground Forums at Scale. In *Proceedings of the 2018 World Wide Web Conference*. ACM, New York, 1845–1854.

[20] Paul A Pavlou. 2003. Consumer acceptance of electronic commerce: Integrating trust and risk with the technology acceptance model. *International Journal of Electronic Commerce* 7, 3 (2003), 101–134.

[21] Kimberly Kiefer Peretti. 2008. Data breaches: What the underground world of carding reveals. *Santa Clara Computer & High Tech. Law Journal* 25 (2008), 375.

[22] Rebecca S. Portnoff, Sadia Afroz, Greg Durrett, Jonathan K. Kummerfeld, Taylor Berg-Kirkpatrick, Damon McCoy, Kirill Levchenko, and Vern Paxson. 2017. Tools for Automated Analysis of Cybercriminal Markets. In *Proceedings of the 26th International Conference on World Wide Web*. ACM, New York, 657–666.

[23] Zhibo Sun, Carlos E. Rubio-Medrano, Ziming Zhao, Tiffany Bao, Adam Doupé, and Gail-Joon Ahn. 2019. Understanding and Predicting Private Interactions in Underground Forums. In *Proceedings of the Ninth ACM Conference on Data and Application Security and Privacy*. ACM, New York, 303–314.

[24] Bruce W Tuckman. 1965. Developmental sequence in small groups. *Psychological Bulletin* 63, 6 (1965), 384.

[25] Rolf van Wegberg, Fieke Miedema, Ugur Akyazi, Arman Noroozian, Bram Klievink, and Michel van Eeten. 2020. Go See a Specialist? Predicting Cybercrime Sales on Online Anonymous Markets from Vendor and Product Characteristics. In *Proceedings of The Web Conference 2020*. ACM, New York, 816–826.

[26] World Health Organisation. 2020. WHO COVID-19 timeline. https://who.int/news-room/detail/29-06-2020-covidtimeline

[27] Michael Yip, Craig Webber, and Nigel Shadbolt. 2013. Trust among cybercriminals? Carding forums, uncertainty and implications for policing. *Policing and Society* 23, 4 (2013), 516–539.

APPENDIX

Contract Process in Detail

Figure 14 shows the contract making process in detail, with 9 transaction statuses denoted by dashed lines. Note for the analyses in this paper, we simplify both 'Complete' and 'Completed' as 'Complete'.

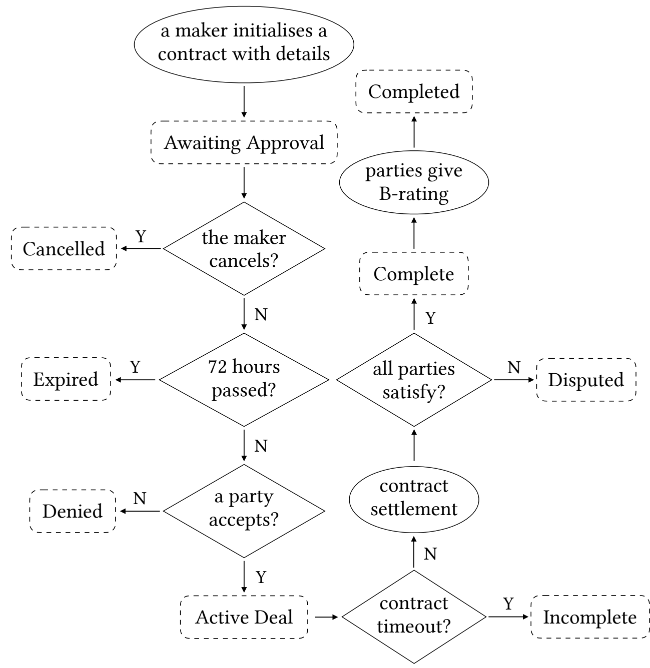

Figure 14: Contract process in detail

Ethical Considerations

This dataset is collected on a publicly available online forum, in which our scraping method does not violate any of the forum's regulations at that time the data was collected. It would be nearly impossible to gain consent from all forum's members as it would be regarded as spamming, however, since our work only studies the collective behaviours instead of individual, and the data are publicly available, it is justified in accordance with the British Society of Criminology's Statement on Ethics [4]. Additionally, in our paper, the identity of all contracts, posts, threads, and users involved are entirely hidden to ensure that no private sensitive information could be revealed. We strictly follow our institution's ethical review procedure and carefully designed our experiments to operate ethically and collectively without aiming to identify individuals. Subject to a strict legal framework, the dataset is provided for academic researchers by the Cambridge Cybercrime Centre so that other researchers can fully reproduce our experiments or conduct further analyses upon our results. In contrast to much previous research on underground cybercrime forums, in this paper we name the site which we investigated as the 'Hack Forums' forum. This is due to the particular characteristics of this forum, which make effective obfuscation impossible: its size (the largest underground forum by some measure) and the operation of its contracts

Table 8: Top 3 transaction flows within each product category over three eras as percentage of all transactions of that kind. E.g., 7% of Exchange during Set-up were made by class F users and accepted by class E users, with an average of 58.8 Exchange between users in these classes per month

	Set-up			Stable			Covid-19		
	Flow	Avg. no. txns month	% transactions of type w/in period	Flow	Avg. no. txns month	% transactions of type w/in period	Flow	Avg. no. txns month	% transactions of type w/in period
Exchange	F → E	58.8	7%	F → K	87.0	7%	F → K	146.0	10%
	F → K	53.1	6%	F → E	60.8	5%	F → E	87.5	6%
	D → B	51.6	6%	G → D	59.2	5%	G → D	78.0	5%
Purchase	H → C	84.8	22%	H → C	106.4	23%	H → C	169.5	26%
	J → C	80.3	20%	J → C	90.2	19%	J → C	121.0	18%
	H → E	25.8	7%	H → K	26.9	6%	H → I	36.5	6%
Sale	C → J	112.1	22%	C → L	967.0	47%	C → L	1123.5	42%
	C → A	67.9	13%	C → A	407.2	20%	C → A	481.5	18%
	I → J	30.8	6%	C → J	191.8	9%	C → J	250.5	9%

Table 9: Zero-Inflated Poisson Regression (All Users of the Contract System)

	Set-up				Stable				Covid-19			
	Estimates		Std. Error	Z Value	Estimates		Std. Error	Z Value	Estimates		Std. Error	Z Value
Count Model Coefficients												
Disputes	0.668	***	0.049	13.504	0.242	***	0.006	42.307	0.107	***	0.011	9.653
Positive Rating	0.173	***	0.007	25.412	0.065	***	0.001	97.093	0.138	***	0.005	27.160
Negative Rating	0.519	***	0.017	30.499	0.084	***	0.009	8.875	0.250	***	0.038	6.644
Marketplace Post Count	0.042	***	0.001	54.019	0.059	***	0.001	101.687	0.057	***	0.001	38.177
No. of Initiated Contracts	0.213	***	0.002	109.568	0.164	***	0.001	186.428	0.267	***	0.002	107.933
No. of Accepted Contracts	-0.042	***	0.006	-7.454	0.019	***	0.000	41.091	0.095	***	0.001	148.330
First-Time Contract Users					-0.780	***	0.009	-83.314	-0.463	***	0.018	-25.455
Length	0.004	***	0.000	23.899	0.000	***	0.000	4.389	0.001	***	0.000	5.267
(Intercept)	0.587	***	0.013	46.734	1.214	***	0.010	120.736	0.575	***	0.015	39.164
Zero-Inflated Model Coefficients												
Disputes	-0.817		0.763	-1.070	-0.345	***	0.104	-3.311	-0.669	***	0.147	-4.540
Negative Rating	-0.549	*	0.228	-2.405	-0.768	**	0.264	-2.910	-0.863		0.770	-1.121
First-Time Contract User					-0.219	***	0.056	-3.888	-0.146	*	0.069	-2.103
Length	-0.009	***	0.001	-7.950	0.000		0.001	-0.427	-0.001		0.001	-0.875
(Intercept)	-0.941	***	0.053	-17.873	-1.316	***	0.059	-22.384	-0.513	***	0.051	-10.029
n	6278				19657				7844			
% of Zero Completed Contracts	27.2				26.6				42.0			
McFadden's R-squared	0.687				0.707				0.654			

*significant at P < 0.05 level; **significant at P < 0.01 level; ***significant at P < 0.001 level.

marketplace mean that it is trivial to identify. As a result, we have made the decision to reject the cover which trivial anonymisation or pseudonymisation would give to this paper (while protecting individual users and transactions). By avoiding the pretense that this forum is not identifiable, we shift the focus to an actual accounting of the potential harms associated with the different kinds of data and analysis which we present.

Coefficients of Zero-Inflated Poisson Models

The coefficients of the models are presented in Tables 9 and 10.

Table 10: Zero-Inflated Poisson Regression (First-Time and Existing Users of the Contract System)

	Stable						Covid-19					
	First-Time Users			Existing Users			First-Time Users			Existing Users		
	Estimates	Std. Error	Z Value	Estimates	Std. Error	Z Value	Estimates	Std. Error	Z Value	Estimates	Std. Error	Z Value
Count Model Coefficients												
Disputes	0.338 ***	0.014	24.240	0.197 ***	0.007	30.040	0.310 ***	0.034	9.251	0.118 ***	0.012	9.935
Positive Rating	0.185 ***	0.005	39.020	0.071 ***	0.001	101.202	-0.314 ***	0.032	-9.704	0.161 ***	0.005	29.612
Negative Rating	-0.820 ***	0.025	-33.440	0.044 ***	0.011	4.031	0.347	0.204	1.701	0.183 ***	0.039	4.725
Marketplace Post Count	0.033 ***	0.002	20.100	0.046 ***	0.001	66.640	-0.028 ***	0.007	-4.031	0.062 ***	0.002	39.407
No. of Initiated Contracts	0.263 ***	0.002	107.470	0.146 ***	0.001	143.932	0.349 ***	0.010	35.072	0.257 ***	0.003	99.084
No. of Accepted Contracts	0.051 ***	0.003	15.560	0.018 ***	0.000	35.916	0.286 ***	0.016	18.422	0.094 ***	0.001	144.777
Length	0.002 ***	0.000	13.120	0.000 **	0.000	2.616	0.000 ***	0.001	-0.831	0.001 ***	0.000	7.081
(Intercept)	-0.133 ***	0.011	-12.570	1.527 ***	0.011	137.967	0.066 *	0.030	2.214	0.541 ***	0.016	33.907
Zero-Inflated Model Coefficients												
Disputes (Set-up)				0.206	0.668	0.308						
Negative Rating (Set-up)				0.312	0.271	1.151						
Disputes (Stable)	0.886 ***	0.210	4.223	-0.850 ***	0.171	-4.973				-0.207	0.132	-1.572
Negative Rating (Stable)	-11.658	102.314	-0.114	-0.337	0.339	-0.992				0.136	0.241	0.562
Disputes (Covid-19)							-0.140	0.228	-0.615	-0.756 ***	0.189	-3.993
Negative Rating (Covid-19)							-13.137	614.755	-0.021	-0.550	0.782	-0.703
Length	0.012	0.002	6.033	-0.004 **	0.001	-3.230	-0.004	0.002	-1.788	0.000	0.001	-0.071
(Intercept)	-3.618 ***	0.185	-19.583	-1.041 ***	0.082	-12.662	-0.674 ***	0.074	-9.149	-0.544 ***	0.056	-9.654
n	16123			3534			2569			5275		
% of Zero Completed Contracts	27.8			20.7			45.8			40.1		
McFadden's R-squared	0.528			0.762			0.505			0.671		

*significant at P < 0.05 level; **significant at P < 0.01 level; ***significant at P < 0.001 level.