

Breaking the Ice: Using Transparency to Overcome the Cold Start Problem in an Underground Market

Tina Marjanov, Konstantinos Ioannidis, Tom Hyndman, Nicolas Seyedzadeh, Alice Hutchings
University of Cambridge
{tm794, ki287, th674, ns896, ah793}@cam.ac.uk

Abstract—New users in anonymous underground marketplaces face the cold start problem; they have no reputation to signal their trustworthiness, hindering their ability to trade in the market. This study explores *transparency*, captured by the choice to make full contract details publicly visible to other users, as an overlooked trust signal. Analysing data from the Hack Forums marketplace, we find that transparency acts as a trust signal, primarily benefiting newcomers constrained by their lack of experience and reputation. However, as participants accumulate credibility, the importance of transparency wanes. Our findings underscore the dynamic nature of trust in such environments, highlighting the distinctive characteristics of private and public contracts as the market evolves.

Index Terms—transparency, trust, signalling, underground markets, cold start problem, cybercrime, asymmetric information

I. INTRODUCTION

Underground forums and cybercrime are intertwined. The forums act as a central hub for information diffusion between cybercriminals, but also provide a marketplace for those involved in cybercrime to trade goods and services ranging from malware to coupon codes for fast food restaurants. Instances of such forum activity include *carding*, which refers to the trade of stolen credit card information [1, 2, 3], *eWhoring*, which refers to an online fraud where users imitate attractive users using pictures, videos, and conversations in exchange for money [4, 5, 6], and *cashing out*, which refers to attempts to legitimise illicit proceeds through currency exchanges [7, 8, 9, 10, 11]. The demand for cybercrime has grown to the point that very limited technical knowledge is needed to commit such crimes as *Cybercrime-as-a-service* has emerged as a profitable business model [12, 13, 14, 15].

Just like legitimate online markets, underground online markets need trust to function. However, in underground markets participants typically do not know who they are trading with and lack concrete evidence about the trustworthiness of other participants. Additionally, they cannot turn to law enforcement or other refund mechanisms in case a trade goes wrong. As such they are vulnerable to *rippers*, i.e., participants who may never fulfil their part of a deal [2]. This uncertainty over the identity of the other agent and the quality of the products being traded has led to problems of asymmetric information plaguing the underground economy, with evidence confirming the suspicion that the market is a *market for lemons* [16, 17].

Operators of underground markets want the markets to succeed. Thus, they have high incentives to provide market-level mechanisms that can be used to signal trust and combat this uncertainty. Campobasso & Allodi [18] provide an overview of market-level mechanisms and their effect on the success of a marketplace. Such mechanisms involve among others: reputation systems where participants rate each other and even leave feedback on the forum after (un)successful trade [19, 3, 20, 21], third-party contract enforcement through a system of administrators and moderators [22, 23], restricted access to (parts of) the market that participants can gain access after verification, payment for upgraded accounts or existing participants vouching for new entrants [17, 20], and a market organised dispute resolution system where participants can attempt to seek justice if they were treated unfairly [24, 20].

How can participants signal their trustworthiness if market-level mechanisms have not been implemented yet or if the existing market-level mechanisms work against them? For instance, while reputation systems may be useful for established actors and help good quality sellers achieve market dominance, new actors may find it difficult to compete as they have no reputation built up yet. The conundrum faced by new actors is known as the *cold start problem*. To overcome the cold start problem, new actors often need to rely on individual-level mechanisms to signal their trustworthiness.

A small but growing literature focuses on documenting individual-level mechanisms to signal trust. Newer participants in such marketplaces aim to overcome the cold start problem by initially making low-value contracts [9], using specialised slang and jargon that other participants associate with experienced participants [25], or providing more details in their advertisements [26]. An under-explored tool in the trust-building toolkit of underground market participants is transparency, defined here as the choice of making contract details publicly visible to all participants of the forum.

For our research, we analyse data from an underground forum with a marketplace where contracts are private by default, but participants sometimes consciously choose to make them public. While the vast majority of contracts in such markets take place behind closed doors with very limited information about each contract being visible, we investigate whether newcomers aim to overcome the cold start problem and increase their perceived trustworthiness by initiating publicly observed contracts.

Specifically, we aim to answer three research questions.

First, we ask whether transparency is used as a trust signal to overcome the cold start problem. Second, we further investigate possible motivations behind the use of transparency. Lastly, we investigate how reliance on transparency as a trust signal changes when market-level mechanisms get established and participants gain credibility. To answer our research questions, we provide a case study using data from Hack Forums. Hack Forums is a prominent and popular underground forum/marketplace, which has the additional feature that market-level mechanisms were gradually established, and our data covers the full period from various stages of the marketplace.

Our key contribution in this paper is documenting that transparency is a valuable individual-level trust signal, which can operate when market-level mechanisms such as reputation or experience are absent or as a substitute for them for unestablished participants. We find that transparency is primarily utilised by (or imposed upon) new participants who are constrained in signalling their trustworthiness by their lack of experience and reputation. Once market-level mechanisms are established and participants gain sufficient experience and reputation, the value of transparency as a trust signal diminishes.

Our summarised findings are:

- Public contracts are used as trust signals primarily when a contract involves at least one new (perceived untrustworthy) participant and mostly before the market is fully established.
- The two motivations behind the use of public contracts are signalling one’s own trustworthiness and as a device for enhanced accountability when trading with a new partner.
- Over time, public contracts (individual-level mechanism) are abandoned and participants rely on experience and reputation (market-level mechanisms).

II. RELATED WORK

Mechanisms for establishing trust in anonymous *legitimate* markets have been extensively studied. Resnick and Zeckhauser [27] find that most users leave feedback on eBay, and the feedback is primarily positive. They also find that feedback is predictive of future activity, which further emphasises the value of reputation systems. Utz *et al.* [28] observe similar patterns studying comments and feedback on Dutch eBay. To further establish the causal relation between reputation and market outcomes, Resnick *et al.* [29] conduct a field experiment where an experienced seller sells identical products using either their established account or a new account, finding that buyers are willing to pay 8% higher prices to buy from the established account.

Mirroring legitimate markets, the majority of *underground* markets offer some type of rating system that allows users to build reputation over time. Hardy and Norgaard [21] studied one of the first darknet markets (the Silk Road) and showed that the reputation system was successful in filtering out bad quality users. Additional to reputation, other mechanisms are found to be effective. Allodi *et al.* [20] compare a failed and

a successful market and found that regulations such as post-trade in-forum exposure of scammers played a key role in the successful market. Yip *et al.* [19, 3] find that a hierarchical structure with administrators, moderators, reviewers, reviewed vendors, and members helped the success of an online carding forum.

Reputation systems, while successful at keeping low quality traders off the market, have their limitations as they may lose informativeness over time. Dupont *et al.* [30] find that average ratings decay over time. This is not driven by highly rated users capitalising on their reputation to run exit scams. On the contrary, when markets become more successful, they attract many new users who start without a reputation, driving average ratings down. As the majority of users who stay longer in such markets are of good quality, it becomes harder to distinguish quality within the active users based on ratings. Dupont *et al.* [30] find more than 80% of ratings are positive, and Motoyama *et al.* [31] finds more than 98%. Since rating a user and providing feedback is often not mandatory, Holt *et al.* [26] find that providing more details about payment mechanisms and the choice of advertisement language can increase the likelihood of receiving feedback.

While experienced participants may suffer from the limitations on the informativeness of reputation systems, a common problem that new participants face is the cold start problem. In an anonymous market, a new participant has no reputation to signal trustworthiness. Motoyama *et al.* [31] provide estimates of the magnitude of the cold start problem. New users need to have more than 50 comments/posts and have sent/received more than 30 private messages before they can start their first trade. Hardy and Norgaard [21] also find that new users need to establish reputation before making their first trade.

To overcome the cold start problem, new participants for whom market-level mechanisms such as reputation are not helpful, need to rely on individual-level mechanisms to signal trustworthiness. Vu *et al.* [9] find that making low-value contracts, which are easier to complete, helps users establish themselves. Hughes *et al.* [25] find that one way new users can overcome this problem is by using argot, a specialised slang and jargon that other users associate with experienced users. We find that transparency -i.e. making contracts public rather than private- also addresses the cold start problem on the individual-level. The success of transparency to signal trustworthiness can be misused by ill-intentioned actors who aim to mimic honest traders. For instance, Vasek & Moore [32] note that Bitcoin High-Yield-Investment-Program scammers “usually post the address in order to signal trustworthiness in the service. Any service that attempts to hide the payment addresses would be viewed with suspicion”.

III. BACKGROUND

A. How asymmetric information can lead to lemon markets

Economists have long thought about how asymmetric information between sellers and buyers affects market outcomes, dating back to Akerlof’s seminal “market for lemons” model [16]. When sellers have better information about the

product they are selling than prospective buyers, an unavoidable feature of underground marketplaces, the uncertainty the buyers are facing can lead to a full collapse of the market for good quality products.

To illustrate, consider a scenario where a seller offers a dataset of stolen credit cards and the buyer cannot distinguish before the trade the quality of the dataset. The dataset could be of high quality (i.e. a *peach*) if the data was stolen recently and the credit cards can still be exploited, or of poor quality (i.e. a *lemon*) if they were acquired a long time ago and most credit cards are now locked. Further assume that: (i) a high quality dataset is worth £4,000 to the seller and £5,000 to the buyer, (ii) a poor quality dataset is worth £1,500 to the seller and £3,000 to the buyer, (iii) 40% of datasets are of high and 60% are of poor quality. In the absence of any information about the quality, the buyer would be willing to pay up to £3,800 which is the average value of a dataset.

For a price higher than £4,000, sellers would be willing to sell both high and poor quality datasets. In this case, the price does not help the buyer distinguish the quality of the dataset, so they would reject the contract as they would be getting a value of £3,800 on average. For a price between £3,000 and £4,000, sellers would only be willing to sell the poor quality datasets. In that case though, the buyer will realise they are getting a bad deal, and will not be willing to pay more than £3,000 for it. Only for prices between £1,500 and £3,000 a trade may occur, since even though the dataset is almost certainly of poor quality, the two parties can negotiate a price that is profitable for both of them. Hence, the only datasets that will be sold would be of poor quality. Contrast the scenario above with the case where both parties have full information about the quality of the dataset offered. Since the buyer values both high and poor quality datasets more than the seller, they would make a profitable trade irrespective of quality.

Asymmetric information results in high quality goods being driven out of the market. This is an undesirable feature for any marketplace that wants to be successful. In ordinary markets such problems are mitigated by legally binding features such as quality guarantees or warranties allowing the buyer to return a low quality product and get refunded. Underground markets typically do not offer such options – with the exception of some markets offering escrow services or dispute resolution systems. Thus, our paper tackles an empirically relevant question and helps us understand how participants in underground markets overcome information asymmetry problems.

B. Contracts and reputation systems on Hack Forums

Hack Forums is a large underground forum boasting a membership of over 5 million individuals, with an archive of over 62 million posts spanning across 6 million threads. Although trades have been occurring informally on the platform since its inception, it was not until June 2018 that the forum administrators introduced, and made mandatory in March 2019, a structured contractual framework enabling users to formalise trades among themselves. While trades may

also occur on other underground forums, Hack Forums has a unique, formal contract system which provides detailed micro-level data.

The image shows a screenshot of the 'Start a New Contract' form on Hack Forums. The form is titled 'Start a New Contract' and includes a disclaimer: 'You are about to initiate a contract with another member. Hack Forums is not responsible for any deals that take place on site. Any deals you conduct on Hack Forums are at your own risk.' The form has several sections: 'User' (Enter the profile link or UID of the user you would like to start a contract with.), 'Thread Link' (Enter the thread link or TID of the thread you are starting a contract for.), 'Contract Timeout' (Enter the number of days this contract must be completed in before it is marked as 'Incomplete'. 14 days), 'Payment Address (Optional)' (Enter your payment address for this contract.), 'Your Position' (Choose the option that best describes your position in the contract. Please Select.), 'Public contract?' (If the above checkbox is checked your contract will be able to be viewed by any user.), and 'Middleman/Escrow (Optional)' (Enter the profile link or UID of the user that will act as middleman/escrow for this contract. They must approve the contract as well.).

Fig. 1. Initiating a contract on Hack Forums

Contracts on Hack Forums are agreements between members that allow them to trade goods and services. Within the contract system, one participant (the *maker*) initiates the contract by either filling in a blank contract form (see first page -before obligations are specified- in Figure 1) or finalising a pre-filled form directly from a thread advertising goods or services. The goods or services involved in the contract are described in the *obligations* fields along with any other terms and conditions. Subsequently, the prospective trading partner of the contract can accept (making them an *accepter*). Upon acceptance, the contract becomes *active* and both parties can complete their part of the agreement. If the accepter takes no action or denies the contract, the contract becomes *expired* or *denied* respectively. Once both sides have completed their obligations, they mark the contract as *completed*. Otherwise, the contract will be left *incomplete*. Finally, if one or both of the participating users believe the other party has not fulfilled their obligations, they can open up a *dispute*. If the administrators believe that there is sufficient evidence that one party has wronged the other, they may take negative action against that account. While this may prevent the offending party from future participation in the marketplace, the wronged party is not refunded in any way.

There are five different types of contracts a participant can choose from when making a contract: *exchange*, *purchase*, *sale*, *trade*, and *vouch copy*. Vouch copies are reviews written by trusted users in exchange for a free version of the product and are frequently used to give an indication of quality to other users. We note that the contract types do not reveal the obligations of the users.

Forum users can see basic information of other users' contracts such as contract type, members involved, date of initiation, and status. When a contract is marked as *public* by its maker, additional details such as maker and accepter obligations, payment information, and contract terms are visible as well. Information such as a user's previous ratings, previous contracts and their disputed contracts are visible on a user's

profile. We emphasise that accepters can only accept or deny a contract as it is, but cannot edit a contract. Thus, the choice to make a contract public is entirely in the maker’s hands.

In addition to the contract system, there are two reputation systems on Hack Forums. First, users participating in contracts rate each other on a system called *B-rating*, using *Booyah* for positive, *Bleh* for neutral, and *Boo* for negative ratings. Beyond the B-rating system there is also a wider *popularity* system of voting that reflects the user’s forum posting activity and quality. Users can add or subtract points from a user’s popularity score, with higher ranking users having a higher weight in their votes.

Before the introduction of the contract system in June 2018, the popularity system was the main way to indicate trustworthiness on Hack Forums. However, since the introduction of the B-rating system, forum administrators have indicated that the popularity system is now primarily a system for sharing community opinions of users rather than a marketplace feature. Moreover, there was a mass reset of popularity ratings with the introduction of the B-rating system in June 2018. In addition to exposing dishonest users as rippers, the administrators may also remove the user’s reputation for acting negatively on the forum, therefore signalling to other users that they are not to be trusted.

In addition to the contract and reputation systems, there is an in-forum currency called *Bytes*, also introduced in June 2018. Bytes can be earned primarily through post activity on the forum, whether that be through users initiating new threads or other active participation in the forum. These Bytes can then be used within the forum for a variety of purposes, such as playing gambling games or to gain account upgrades. As Bytes can be donated to other users, this alternative currency is occasionally bought or sold between users.

IV. METHODOLOGY

A. CrimeBB dataset

The CrimeBB [8] dataset contains over 110 million posts and related artefacts scraped from underground forums related to various cybercriminal activities. The dataset is maintained by the Cambridge Cybercrime Centre¹ and shared with academic researchers upon formal agreement. Topics discussed across the forums include hacking, coding, legal and illegal money making methods, malware, trade of different datasets, cheating in online games etc. The forums generally contain a number of topic-specific boards on which users can start threads and make posts. Specific to Hack Forums is the contract system, described in subsection III-B. We use the subset of CrimeBB that includes posts and contracts collected from Hack Forums for our analysis.

In total, the analysed subset contains almost 350,000 unique contracts on Hack Forums, ranging from June 2018 until November 2022. We do not have full data on certain contracts from September 2020 to June 2022 (approximately 120,000 or 35% of all contracts). Specifically, contracts made during

that time period all appear private, even if they may have been originally set to public due to the forum setting all contracts to private six months after a contract has been made. We use the contracts where we do know their true visibility to build a machine learning classifier that predicts the original visibility for the contracts in that period (see subsection IV-B). We excluded 4,418 contracts that were still active at the end of the scraping period as their status was not yet determined. The purged dataset includes 343,921 contracts which were created from June 2018 to November 2022.

We notice that the use of sale and purchase type labels is used inconsistently by the forum users. Specifically, we conjecture that most accepters are acting as sellers, even when the contracts are marked as sales by the maker (implying that the maker is the one selling). We provide three pieces of evidence supporting our conjecture. First, twice as many accepters make threads advertising goods or services for sale compared to the number of makers doing so. Second, within contracts initially marked as sales, the accepters are thanking makers for their purchase eight times more often than the other way around. Finally, we confirm the inconsistencies by manually inspecting a random subset of contracts where the feedback included words related to selling or purchase and indeed find that the majority of accepters act as sellers regardless of the contract type. We therefore disregard the original labels as unreliable and consider makers as buyers and accepters as sellers. For clarity, we will use terms buyers and sellers solely in the remainder of the paper.

Due to conceptual similarities and very low frequencies, we also merge contracts that were invalid, expired, denied, and cancelled with incomplete contracts for the remaining of the paper. This leaves us with three contract statuses: complete contracts, incomplete contracts, and disputed contracts.

B. Public contract classifier

We design a binary classifier to label the public contracts that were incorrectly collected as private. The classifier uses contract characteristics (e.g., contract type, completion status, date) as well as additional user characteristics as features (e.g., total contracts, reputation at time of contract, reputation currently, previous amount and proportion of public contracts). We use XGBoost [33], a tree boosting algorithm that has performed well in the past in many machine learning tasks with imbalanced datasets where others have struggled with overfitting [34, 11, 35]. We perform Grid Search Cross-Validation to select the hyperparameters: $\gamma=1$, $\text{learning_rate}=0.3$, $\text{max_depth}=5$, $\text{reg_lambda}=10$ with the objective being logistic regression for binary classification that outputs probability. This gave a final 10-fold cross validation scores of: 93% accuracy, 74% precision and 55% recall. We use the classifier to complete the gaps in the dataset that occurred due to forum’s policy changes.

C. Ethics

We obtained approval from the Department’s research ethics committee. CrimeBB data is collected through the use of web

¹<https://cambridgecybercrime.uk/>

scrapers, and hence informed consent has not been obtained by the individuals concerned. Established ethical guidelines for online research into criminal activity recommend that informed consent may not be required for research into online communities where the data is publicly available and the research outputs focus on collective rather than individual behaviour [36]. To further reduce the likelihood of harm to any users of these platforms as a result of the research, we do not publish identifying information, and handle the collected data securely and sensitively.

V. RESULTS

A. Marketplace overview and user characteristics in Hack Forums

Between June 2018 and December 2022, roughly 350,000 unique contracts have been recorded on the Hack Forums marketplace. Figure 2 shows the total number of contracts started across that period as well as the proportion of the contracts that were successfully completed. There is a distinct bump in the number of contracts started after March 2019. The bump coincides with Hack Forums making the contracts mandatory. We also see that the ratio of completed contracts drops heavily around the same time.




Fig. 2. Total and complete contracts over market lifespan

The vast majority of contracts are private (95.14%) with only a small portion being public. Most contracts are sales and purchases of products (78.59%) followed by exchanges (19.45%). Trades and vouch copies appear in only 0.91% and 1.05% of the contracts respectively. Most contracts were either completed (40.68%) or incomplete (58.37%), with a very small number of contracts being disputed (0.95%). While the majority of the contracts (70.82%) were arranged between users with no previous interaction, a substantial number of users made multiple contracts with each other.

To gather some insight into what is being traded, a feature of a contract that users engaging in criminal trades may want to obfuscate, we qualitatively analyse a subset of public contracts. We manually classified the goods or services of 2,500 public contracts into 6 distinct groups: exchange of cryptocurrencies (30%), non-explicitly malicious services and tools (22%), trade of accounts (20%), explicitly malicious services or tools

(12%), in-forum reputation trading (10%) and video game-related trades (3%).

The market consists of 71,105 unique users among which 55,604 only appear as buyers, 3,787 only appear as sellers, and 11,714 have participated in contracts in both roles. Activity of buyers and sellers is vastly different as the market is characterised by very active sellers who participate on average in 35 contracts and less active buyers who participate on average in 10 contracts. Around half (50.57%) of unique buyers and 16.8% of unique sellers only engaged in a single contract. On the other hand, the top 5% of buyers engaged in at least 26 contracts each and the top 5% of sellers in at least 99 contracts each.

Our observations about the activity and the roles of users indicate that the market is highly centralised, with a much smaller number of sellers supplying a much larger number of buyers. Buyers only rarely participate in the marketplace as sellers, whereas sellers will occasionally buy. The overall pattern suggests parallels to the way legitimate online markets work. A smaller number of sellers operate as suppliers of the marketplace to a much larger number of buyers who buy their goods and services. To illustrate, the seller to buyer ratio on eBay was roughly 1 to 8 in 2022.²

B. Transparency as a trust signal

In this section we address our first research question. Specifically, we aim to show that transparency of a contract makes it more likely for the contract to be completed, and that transparency as a signal is used primarily by unestablished users.




Fig. 3. Visibility across contract status

Figure 3 provides bar graphs of the transparency of contracts across contract status. More specifically, we present the two conditional distributions of contract status, namely for public contracts (left bar within each contract status) and for private contracts (right bar within each contract status). We observe two patterns: (i) 49.5% of public contracts are completed whereas only 40.2% of the private contracts are, and (ii) 58.9% of private contracts are incomplete whereas only 47.6% of the

²According to the 2022 annual report [37], eBay had 138 million people are active on eBay, out of which only 18.3 million were sellers.

public contracts. Formal χ^2 tests verify that contract status does depend on transparency of a contract ($p\text{-value} < 0.001$). Public contracts have higher completion rates compared to private contracts. Hence, there is scope for users to use transparency to their advantage.

We proceed by investigating who and when uses transparency by looking at which point in their lifespan the users make public contracts. Figure 4 plots the number of public contracts over the lifespan of a user. We emphasise that in the graph, we standardise the lifespan of all users in the same time horizon, so that for every user the leftmost part of the graph indicates their first contract and the rightmost part of the graph indicates their last contract. For instance, for users with few contracts, their last contract may be their third, whereas for very active users their third contract may still be very early in their lifespan.




Fig. 4. Public contracts during the lifespan of buyers and sellers

Among the buyers, most will only make a single public contract in their marketplace lifespan (45.6%), 18.4% will make only two and 10.3% will make three. Among sellers, the pattern is similar, with 40.4%, 18.3%, 10.2% accepting only one, two, and three public contracts, respectively. Among both buyers and sellers less than 5% made more than 15 public contracts. The user with the most public contracts has made 227, appearing in both roles.

We find a distinct spike in public contracts at the beginning of a user's lifespan. We see that the public contracts disproportionately involve new users, who do not yet have any reputation within the marketplace. The users tend to quickly switch to private contracts in the future with the majority of them doing so already on their second contract.

C. The motivations for using transparency

Having established that transparency -in the form of engaging in a public contract- is used strategically by participants in the marketplace, we proceed with our second research question. We are interested in investigating the motivations behind the use of transparency. Two plausible motivations of participants to engage in a public contract are: (i) signalling their own trustworthiness and (ii) ensuring the trustworthiness of their trading partner to minimise the chance of getting

scammed. The first motivation primarily applies to participants who themselves are inexperienced, whereas the latter primarily applies when the trading partner is inexperienced.

To continue our analysis, we need to classify participants into experienced and inexperienced. As mentioned earlier in subsection V-A, approximately half the participants only participated in a single contract. Hence, we classify a participant as inexperienced when they have participated in at most two contracts, and as experienced otherwise.³ Thus, we generate four experience groups as the combinations of experience levels and market roles. Table I shows the distribution of experience in the first two years of the marketplace.

TABLE I
DISTRIBUTION OF EXPERIENCE GROUPS

Experience		Market year	
Buyer	Seller	First	Second
Low	Low	1,033 (1.66%)	543 (0.46%)
High	High	38,856 (62.60%)	83,189 (71.07%)
High	Low	4,708 (7.58%)	3,593 (3.07%)
Low	High	17,478 (28.16%)	29,789 (25.40%)
Total		62,075 (100.00%)	117,054 (100.00%)

Column percentages in parentheses.

Transparency takes different forms for buyers and sellers. Buyers choose whether to initiate a contract as public or private, whereas sellers choose whether to accept or reject a contract for which the visibility status is already determined. We therefore investigate the presence of proposed motivations for buyers and sellers separately.

We first investigate the buyers' decisions. We find evidence for both motivations. Figure 5 plots the ratio of public contracts initiated by buyers across experience groups.

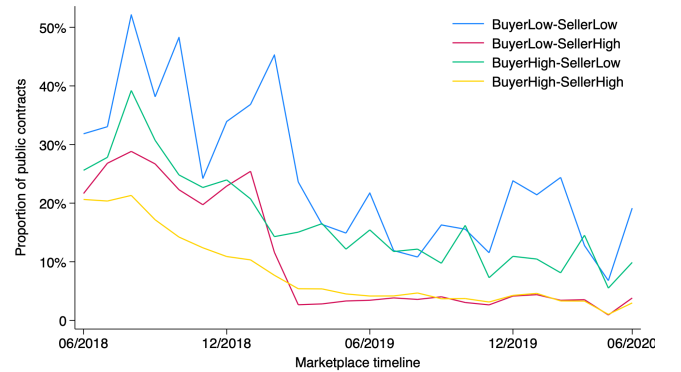


Fig. 5. Ratio of public contracts across experience groups

When interacting with inexperienced sellers, inexperienced buyers choose public contracts more often than experienced buyers. This is evident from the fact that the blue line in the figure is above the green line. When facing experienced sellers, we observe a similar pattern as the red line is above

³The results presented in this subsection are robust to alternative experience thresholds.

the yellow line, especially during the first year. Those two patterns provide evidence that new buyers are motivated by the desire to signal their own trustworthiness.

Fixing the experience level of the buyer, we can look at whether buyers are more likely to request public contracts when trading with inexperienced sellers. We find evidence that this is the case both for inexperienced buyers (blue line is above the red line) and for experienced buyers (green line is above the yellow line). While those patterns provide clear evidence that buyers want to impose higher accountability when facing inexperienced sellers, we stress that they rarely have the need to do so as the market is dominated by experienced sellers and contracts with inexperienced sellers are rare - approximately 9%.

Next, we look at the sellers' motivations. Given that the publicity of the contract is determined by the buyer who typically initiates the contract, we look at the choice of sellers to accept or reject contracts. We separately compute the acceptance rates of public contracts and private contracts over time, and in Figure 6 plot the difference in acceptance rates between public and private contracts. When a line is above zero, this indicates that the acceptance rate for public contracts is higher than the acceptance rate of private contracts, and vice versa.




Fig. 6. Difference in acceptance rates between public and private contracts

Inexperienced sellers accept public contracts more often than experienced sellers both when the contract is initiated by an inexperienced buyer (blue line is above zero) and when the contract is initiated by an experienced buyer (green line is above zero). Thus, we find evidence that inexperienced sellers do want to signal their own trustworthiness by accepting public contracts more often than private contracts.

We find weaker evidence for the second motivation of sellers demanding accountability from new buyers. Inexperienced sellers do accept public contracts more often when facing inexperienced buyers than when facing experienced buyers (blue line is above green line), suggesting that inexperienced sellers do request accountability from new buyers to some extent. The evidence is unclear for experienced sellers as they only demand accountability from buyers very early on in the marketplace (red line above zero initially), but very quickly

they switch to strongly prefer private contracts (both red and yellow lines are below zero).

We end this section with a comment on the interaction of the motivations between buyers and sellers. In both figures we observe that after the first year, buyer experience does not affect the ratios of public contracts or the acceptance rate of public contracts as much - see that the blue line is close to the green line, and the red line is close to the yellow line in both figures. On the contrary, only the experience of the seller matters, with inexperienced sellers participating in public contracts more frequently than experienced sellers. This means that buyers still demand public accountability from the rarely occurring inexperienced sellers. We also note this is the time when the market stabilises with a number of experienced power-sellers running the marketplace. This, in turn makes the inexperienced sellers the rare oddity, which might explain the extra caution exercised by the buyers.

However, this also means that after controlling for seller experience, there is no difference in public contracts between experienced and inexperienced buyers. This can be driven by either reluctance of buyers to initiate public contracts against experienced sellers or by the reduced willingness of experienced sellers to accept public contracts. We find evidence for the latter by looking more closely at the proportions of cancelled, denied or expired contracts. For experienced sellers, the proportion of those contracts when the buyer initiated a public contracts is 27.7%, but only 11.5% when private. On the contrary, inexperienced buyers cancel, deny or let expire 22.6% of public contracts and 28.7% of private. This suggests that experienced sellers can implicitly enforce the use of private contract, possibly because they can afford the extra risk of trading with an inexperienced buyer. Inexperienced buyers only have the option to signal their own trustworthiness through the use of public contracts when facing an inexperienced seller.

D. The increased reliance on credibility over transparency

Having documented that transparency is indeed used, by whom and why, we proceed to investigate whether, and if so how, reliance on transparency changes when the marketplace matures and market-level mechanisms get established. We already observed that experienced sellers' tolerance for transparency drops after some time. We now also explore what other trust signals take over in its stead.

Figure 7 plots the proportions of public and first contracts over the lifetime of the marketplace. We see that in the beginning of the marketplace there are over 20% public contracts, but this slowly decays to around 5% after a few months of the market being active. One plausible reason for the low propensity of public contracts could be that there are not many new participants who would need it in the first place. However, the proportion of first contracts also starts high at around 40% and quickly stabilises around 20% which indicates that the marketplace does attract new participants, but those new participants rely on public contracts less over time.




Fig. 7. Public and first contracts during the lifespan of the marketplace

An alternative reason is that market-level mechanisms substitute the usefulness of transparency as a trust signal. To check for this explanation, we perform a mediation analysis.⁴ Our analysis proceeds in three steps: (i) verify econometrically that indeed visibility of a contract significantly increases the probability that a contract is completed, (ii) find variables that can explain the choice of whether a contract is public or private, and (iii) re-estimate the effect of visibility on completion rates after controlling for those variables. In our case, we investigate two mediator variables (market-level trust signals), namely experience and reputation. We measure experience by the number of contracts a participant has engaged in at the time of a given contract and reputation by the number of completed contracts at the time of a given contract.

Variable name	(1) Status	(2) Visibility	(3) Status	(4) Status	(5) Visibility	(6) Status
Visibility (public)	0.279*** (0.024)		0.003 (0.024)	0.279*** (0.024)		0.422*** (0.022)
Buyer experience		-0.006*** (0.001)	0.002*** (0.000)			
Seller experience		-0.004*** (0.000)	-0.001*** (0.000)			
Buyer reputation					-0.006*** (0.001)	0.006*** (0.000)
Seller reputation					-0.006*** (0.000)	0.000*** (0.000)
Observations	203,703	203,703	203,703	203,703	203,703	203,703

Each column is a logistic regression with dependent variable in the column title. Standard errors clustered on user ID (in parenthesis). Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Includes observations up to September 2020.

TABLE II
EXPERIENCE AND REPUTATION ON VISIBILITY AND COMPLETION

Table II presents results from logistic regressions. Given that the observations in our dataset -the contracts- are not independent since many users have engaged in multiple contracts, we cluster errors on the user level. We only use observations until September 2020 since public contracts are too few for meaningful analysis later on.

We first estimate the effect of visibility of a contract on completion rates and verify econometrically that indeed public contracts are significantly more likely to be completed

⁴This method is often used in psychology to document how the effect of an independent variable on a dependent variable is mediated through an intermediate *mediator* variable.

(column 1). Next, we document that experience is strongly correlated with the choice of whether a contract is public or private with more experienced participants choosing private contracts (column 2). Finally, we show that after controlling for experience, the effect of visibility on completion rates disappears as it is no longer significant (column 3). Columns 4-6 show that controlling for reputation instead of experience does not eliminate the effect of visibility on completion rates.⁵

Our results complement the observation from Figure 4 that more experienced users choose to engage in private contracts -in other words experience predicts visibility-, but conditional on that choice, the visibility of the contract on its own has no additional effect on the likelihood that a contract is completed.

The fact that more experienced users also have more completed contracts implies that the informational value of reputation is to a large extent captured by the experience of a user. As also found in other studies ([31, 30]), when the vast majority of ratings are positive it is harder to distinguish trustworthiness of users based on reputation. Thus, it is not surprising that controlling for reputation does not mediate the effect of visibility, whereas controlling for experience fully mediates the effect. The overall pattern roughly suggests that credibility as measured by either experience or reputation of participants is the key to understanding which contracts are completed in established markets, whereas transparency plays an important role in early markets and limited later on.

E. Qualitative analysis of major forum threads

In order to grasp the participants' perspectives, perceptions, and motivations, we supplement our findings with a qualitative analysis of posts made by forum administrators, moderators and participants of the contract system. We limit the analysis to the major threads dedicated to the introduction and the explanation of the contract system, major announcements (e.g. making contracts mandatory), advice (e.g. how to not get scammed) and most popular discussions related to the contract system and its individual components.

The forum administrators are quite open about the great effort they went through in order to establish Hack Forums as a trustworthy and reliable marketplace. The marketplace went through several iterations before the contract system was established. The administrator states that before the contract system the marketplace was "like the Wild West and protecting yourself from scammers was something you had to do on your own". The sentiment among participants is similarly positive when talking about the contract system. They almost uniformly agree that the usage of contracts leads to a safer and more trustworthy marketplace. However, there is also a widespread belief that trust on the market is only possible if the adoption of contracts is universal.

Before contracts were mandatory, users could make transactions without a contract. We could view the choice to make a private contract as an instance of partial transparency and

⁵Given that the vast majority of contracts that are being rated are completed contracts, using average rating instead of number of complete contracts produced qualitatively similar results.

the choice to make a public contract as full transparency. This suggests that the high completion rate before March 2019 as observed in Figure 2 is driven by the self-selection of more trustworthy users into making contracts. In March 2019 the contracts are eventually made mandatory. Again, the move is praised with users expressing beliefs that scams are unlikely to happen, provided that users utilise the contract system correctly. This corroborates our findings suggesting that market-level mechanisms are successful in establishing overall trust. Specifically, contracts becoming mandatory coincides with a drop in transparency usage, especially when big vendors are involved.

However, we also find posts made by experienced users and forum administrator outlining the individual steps that users should take in order to protect themselves. We find strong emphasis on knowing who you're trading with and "doing your homework". Users explain that some of the steps they take include reviewing the partner's trading profile and their previous contracts, checking whether the partner is a part of a trusted group, whether the partner has donated to the marketplace, whether their previous posts and contracts are consistent with a specific business venture etc. Not surprisingly, new users are treated with more suspicion.

This is interesting for two reasons. First, we can see that individual-level mechanisms to establish trust do appear and are relied upon even in established markets. Second, getting a glimpse at experienced user's thoughts allows us to better understand new user's motivations for their actions. If new users understand that their potential trading partners will be examining their profile before engaging in a trade with them, it is in their best interest to be transparent about their previous actions in order to send a positive signal about their trustworthiness.

VI. CONCLUDING DISCUSSION

In this paper, we document the use of transparency by unestablished participants on underground marketplaces. The need for transparency arises due to new users' inability to rely on market-level trust mechanisms as they have no experience or reputation built up yet. Instead, new users need to find individual-level ways to signal trust to perspective trading parties. This is not only beneficial for the success of the present contract, but it also gives prospective trading partners a better insight into the user's market history.

We find that the reliance on transparency is stronger in the early months of the marketplace when the majority of users as well as the marketplace itself were still finding their footing. Around the one year mark we find a significant drop in the usage of public contracts, especially among the (now) established users. We stress that the relative drop in public contracts coincides with (and is likely a result of) the decision of the marketplace to make the contracts mandatory for everyone involved in trading. This likely led to an increase in contracts made by traders that were less enthusiastic in using the contract system in the first place, and therefore unlikely to make them public.

We emphasise the significance of examining individual-level trust signals. This is especially relevant for emerging anonymous marketplaces which may lack fully established market-level mechanisms. Furthermore, even in the presence of such mechanisms, the continuous growth of the market introduces new participants who may not have accumulated sufficient reputation to serve as a reliable trust signal. Therefore, the insights drawn from our case study on Hack Forums contribute not only to our understanding of this specific platform, but also to the broader consideration of trust dynamics in evolving and expanding online marketplaces. Our study also provides insight in the interplay of market-level and individual-level trust mechanisms as we document that experience and reputation (market-level) gradually take over and fully substitute transparency (individual-level).

Participants in underground markets want to hide what they are trading as it often constitutes a crime. Under this assumption, private contracts are the default in such marketplaces, and users only choose to trade in public early on until they establish themselves. What exactly is being traded is of immense importance both for researchers to measure harm being caused, but also for law enforcement to pin down the exact nature of cybercrimes. Typically such underground marketplaces attract attention from researchers and policymakers/law enforcement after they establish themselves as notorious and cause substantive harm. Researchers can obtain better quality data as more information about the activities taking place in the forum may be accessible early on, and even use early data to train classifiers and predict what is being traded in private contracts [38]. Given that various interventions by law enforcement may have different effectiveness in disrupting emerging marketplaces compared to established ones, law enforcement can use visibility of contracts or equivalent mechanisms as a signal of whether a marketplace has matured, and choose their interventions accordingly.

We end our paper by discussing limitations and suggest future research that can address them. Our results are based on a single underground forum and marketplace (Hack Forums) so we can only confidently make claims about that specific case. We hope to have uncovered a promising trust signal with broader applications, so we encourage future research to study multiple marketplaces and provide further evidence on whether indeed transparency is used as an individual-level signal of trust in the absence of market-level signals such as reputation and experience.

By the nature of using field data, our results are correlational and cannot establish causality. Despite partially trading away external validity, experimental approaches could help us better understand the causal links between transparency and credibility in anonymous marketplaces. For instance, Sebahg et al. [39] provide an experiment to measure the efficiency of an intervention aiming to disrupt a marketplace by attacking participant reputation. We believe an experiment which incorporates as many ingredients as possible from underground markets and across treatments varies whether contracts are

always public, always private, or users are allowed to endogenously choose contract visibility could help us make causal claims on how transparency is used as a trust signal and how it affects market outcomes.

ACKNOWLEDGMENTS

We thank the Cambridge Cybercrime Centre for access to the CrimeBB dataset, as well as our colleagues for their support and advice. This work is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 949127). Tina Marjanov is supported by King's College, Cambridge and the Cambridge Trust. Konstantinos Ioannidis gratefully acknowledges support through the Leverhulme International Professorship in Neuroeconomics awarded to the University of Cambridge. Tom Hyndman and Nicolas Seyedzadeh were supported by the King's College Summer Research Programme.

REFERENCES

- [1] J. Franklin, A. Perrig, V. Paxson, and S. Savage, "An inquiry into the nature and causes of the wealth of internet miscreants," in *Proceedings of the 14th ACM Conference on Computer and Communications Security*, vol. 7. ACM, 2007, pp. 375–388.
- [2] T. J. Holt and E. Lampke, "Exploring stolen data markets online: Products and market forces," *Criminal Justice Studies*, vol. 23, no. 1, pp. 33–50, 2010.
- [3] M. Yip, C. Webber, and N. Shadbolt, "Trust among cybercriminals? carding forums, uncertainty and implications for policing," *Policing and Society*, vol. 23, no. 4, pp. 516–539, 2013.
- [4] A. Hutchings and S. Pastrana, "Understanding ewhoring," in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 201–214.
- [5] S. Pastrana, A. Hutchings, D. Thomas, and J. Tapiador, "Measuring ewhoring," in *Proceedings of the Internet Measurement Conference*, 2019, pp. 463–477.
- [6] M. Campobasso and L. Allodi, "Impersonation-as-a-service: Characterizing the emerging criminal infrastructure for user impersonation at scale," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020, pp. 1665–1680.
- [7] R. S. Portnoff, S. Afroz, G. Durrett, J. K. Kummerfeld, T. Berg-Kirkpatrick, D. McCoy, K. Levchenko, and V. Paxson, "Tools for automated analysis of cybercriminal markets," in *Proceedings of the 26th international conference on world wide web*, 2017, pp. 657–666.
- [8] S. Pastrana, D. R. Thomas, A. Hutchings, and R. Clayton, "Crimebb: Enabling cybercrime research on underground forums at scale," in *Proceedings of the 2018 World Wide Web Conference*, 2018, pp. 1845–1854.
- [9] A. V. Vu, J. Hughes, I. Pete, B. Collier, Y. T. Chua, I. Shumailov, and A. Hutchings, "Turning up the dial: the evolution of a cybercrime market through SET-UP, STABLE, and COVID-19 eras," in *Proceedings of the ACM Internet Measurement Conference*. ACM, 2020, pp. 551–566.
- [10] A. Bermudez-Villalva and G. Stringhini, "The shady economy: Understanding the difference in trading activity from underground forums in different layers of the web," in *2021 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 2021, pp. 1–10.
- [11] G. A. Siu, B. Collier, and A. Hutchings, "Follow the money: The relationship between currency exchange and illicit behaviour in an underground forum," in *2021 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2021, pp. 191–201.
- [12] A. K. Sood and R. J. Enbody, "Crimeware-as-a-service—a survey of commoditized crimeware in the underground market," *International Journal of Critical Infrastructure protection*, vol. 6, no. 1, pp. 28–38, 2013.
- [13] A. Mikhaylov and R. Frank, "Cards, money and two hacking forums: An analysis of online money laundering schemes," in *2016 European intelligence and security informatics conference (EISIC)*. IEEE, 2016, pp. 80–83.
- [14] U. Akyazi, M. van Eeten, and C. H. Gañán, "Measuring cybercrime as a service (CaaS) offerings in a cybercrime forum," in *Workshop on the Economics of Information Security*, 2021, pp. 1–15.
- [15] Z. Sun, A. Oest, P. Zhang, C. Rubio-Medrano, T. Bao, R. Wang, Z. Zhao, Y. Shoshitaishvili, A. Doupe, G.-J. Ahn *et al.*, "Having your cake and eating it: An analysis of {Concession-Abuse-as-a-Service}," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 4169–4186.
- [16] G. Akerlof, "The market for lemons: Quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, vol. 84, no. 3, pp. 488–500, 1970.
- [17] C. Herley and D. Florêncio, "Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy," in *Economics of information security and privacy*, 2010, pp. 33–53.
- [18] M. Campobasso, R. Radulescu, S. Brons, and L. Allodi, "You can tell a cybercriminal by the company they keep: A framework to infer the relevance of underground communities to the threat landscape," in *The 22nd Workshop on the Economics of Information Security (WEIS'23)*, 2023.
- [19] M. Yip, N. Shadbolt, and C. Webber, "Why forums? an empirical analysis into the facilitating factors of carding forums," in *Proceedings of the 5th annual ACM web science conference*. ACM, 2013, pp. 453–462.
- [20] L. Allodi, M. Corradin, and F. Massacci, "Then and now: On the maturity of the cybercrime markets the lesson that black-hat marketers learned," *IEEE Transactions on Emerging Topics in Computing*, vol. 4, no. 1, pp. 35–46, 2015.
- [21] R. A. Hardy and J. R. Norgaard, "Reputation in the internet black market: an empirical and theoretical analysis of the deep web," *Journal of Institutional Economics*, vol. 12, no. 3, pp. 515–539, 2016.

- [22] J. Lusthaus, "Trust in the world of cybercrime," *Global crime*, vol. 13, no. 2, pp. 71–94, 2012.
- [23] —, "How organised is organised cybercrime?" *Global Crime*, vol. 14, no. 1, pp. 52–60, 2013.
- [24] D. Georgoulas, J. M. Pedersen, M. Falch, and E. Vasilomanolakis, "A qualitative mapping of darkweb marketplaces," in *2021 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 2021, pp. 1–15.
- [25] J. Hughes, A. Caines, and A. Hutchings, "Argot as a trust signal: Slang, jargon & reputation on a large cybercrime forum," in *Workshop on the Economics of Information Security*, 2023.
- [26] T. J. Holt, O. Smirnova, and A. Hutchings, "Examining signals of trust in criminal markets online," *Journal of Cybersecurity*, vol. 2, no. 2, pp. 137–145, 2016.
- [27] P. Resnick and R. Zeckhauser, "Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system," in *The Economics of the Internet and E-commerce*. Emerald Group Publishing Limited, 2002, pp. 127–157.
- [28] S. Utz, U. Matzat, and C. Snijders, "On-line reputation systems: The effects of feedback comments and reactions on building and rebuilding trust in on-line auctions," *International Journal of Electronic Commerce*, vol. 13, no. 3, pp. 95–118, 2009.
- [29] P. Resnick, R. Zeckhauser, J. Swanson, and K. Lockwood, "The value of reputation on ebay: A controlled experiment," *Experimental economics*, vol. 9, pp. 79–101, 2006.
- [30] B. Dupont, A.-M. Côté, C. Savine, and D. Décary-Héту, "The ecology of trust among hackers," *Global Crime*, vol. 17, no. 2, pp. 129–151, 2016.
- [31] M. Motoyama, D. McCoy, K. Levchenko, S. Savage, and G. M. Voelker, "An analysis of underground forums," in *Proceedings of the 2011 ACM SIGCOMM Conference on Internet measurement conference*, 2011, pp. 71–80.
- [32] M. Vasek and T. Moore, "There's no free lunch, even using bitcoin: Tracking the popularity and profits of virtual currency scams," in *Financial Cryptography and Data Security: 19th International Conference, FC 2015, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers 19*. Springer, 2015, pp. 44–61.
- [33] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 2016, pp. 785–794.
- [34] A. Caines, S. Pastrana, A. Hutchings, and P. J. Buttery, "Automatically identifying the function and intent of posts in underground forums," *Crime Science*, vol. 7, no. 1, p. 19, 2018.
- [35] G. Atondo Siu and A. Hutchings, "'Get a higher return on your savings!': Comparing adverts for cryptocurrency investment scams across platforms," in *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2023, pp. 158–169.
- [36] British Society of Criminology, "Statement of Ethics," 2015.
- [37] eBay, *2022 Annual Report*, 2022, https://ebay.q4cdn.com/610426115/files/doc_financials/2022/ar/2022-annual-report.pdf, accessed on March 28, 2024.
- [38] Y. Fan, Y. Ye, Q. Peng, J. Zhang, Y. Zhang, X. Xiao, C. Shi, Q. Xiong, F. Shao, and L. Zhao, "Metagraph aggregated heterogeneous graph neural network for illicit traded product identification in underground market," in *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2020, pp. 132–141.
- [39] L. Sebahg, J. Lusthaus, E. Gallo, F. Varese, and S. Sirur, "Cooperation and distrust in extra-legal networks: a research note on the experimental study of marketplace disruption," *Global Crime*, vol. 23, no. 3, pp. 259–283, 2022.