



Open Market or Ghost Town? The Curious Case of OpenBazaar

James E. Arps and Nicolas Christin^(✉)

CyLab, Carnegie Mellon University, Pittsburgh, USA
{jarps,nicolasc}@cmu.edu

Abstract. OpenBazaar, a decentralized electronic commerce marketplace, has received significant attention since its development was first announced in early 2014. Using multiple daily crawls of the OpenBazaar network over approximately 14 months (June 25, 2018–September 3, 2019), we measure its evolution over time. We observed 6,651 unique participants overall, including 980 who used Tor at one point or another. More than half of all users (3,521) were only observed on a single day or less, and, on average, only approximately 80 users are simultaneously active on a given day. As a result, economic activity is, unsurprisingly, much smaller than on centralized anonymous marketplaces. Furthermore, while a majority of the 24,379 distinct items listed seem to be legal offerings, a majority of the measurable economic activity appears to be related to illicit products. We also discover that vendors are not always using prudent security practices, which makes a strong case for imposing secure defaults. We conclude that OpenBazaar, so far, has not gained much traction to usher in the new era of decentralized, private, and legitimate electronic commerce it was promising. This could be due to a lack of user demand for decentralized marketplaces, lack of integration of private features, or other factors, such as a higher learning curve for users compared to centralized alternatives.

Keywords: Measurement · Peer-to-peer systems · Electronic commerce

1 Introduction

OpenBazaar, originally called DarkMarket, is a peer-to-peer electronic commerce platform that has attracted significant media attention [10, 13, 14]. It was first envisioned in 2014 as a response to the government takedown of the Silk Road [6] online anonymous (“darknet”) marketplace [10]. However, OpenBazaar developers quickly pivoted away from the darknet marketplace space and toward a decentralized e-commerce platform. The project raised several million dollars in seed funding through a startup called OB1 [13, 14].

While the default OpenBazaar *search engine* is developed by OB1 and filters undesirable items such as narcotics [15], the OpenBazaar project itself is open-source and the developers cannot prevent vendors from listing such items using the platform.

Interestingly, the takedown of the Silk Road marketplace failed to curb illicit activity on centralized anonymous marketplaces. On the contrary, Silk Road’s successors have been thriving [9, 18, 21, 30]: As of 2017, the leading marketplaces were grossing hundreds of millions of dollars per year [9], with no decline in sight, despite adversarial events such as law enforcement operations and “exit scams,” in which some marketplace operators abruptly shut down their servers, taking with them any money left in escrow on the platform.

Given that OpenBazaar, thanks to its decentralized architecture, could prevent or mitigate most of these adversarial events, and given the clear economic demand for such services, we would expect thriving economic activity on a peer-to-peer platform such as OpenBazaar. This is what we measure in this paper.

For approximately 14 months (June 25, 2018–September 3, 2019), we have been crawling the OpenBazaar network on a near-daily basis to get a sense of the number of participants and the size of the overall network. By scraping listings offered by each vendor and the associated feedback left by buyers (similar to techniques used in related work, e.g., [6, 18, 30]), we also estimate the economic activity taking place on OpenBazaar. Finally, we can examine the extent of users’ mindfulness about operational security when trading in illicit or illegal products.

Our findings are rather sobering. While there is undeniable network activity, and a reasonably large number of items available on the OpenBazaar network (at least 24,379 distinct listings observed during our measurement period), economic activity remains modest, and appears to be mostly generated by illicit product sales. Furthermore, many vendors appear to misunderstand the security guarantees offered by OpenBazaar or fall prey to some configuration defaults, and publicly reveal some potentially compromising information about themselves: Nearly a fifth of the vendors that use OpenBazaar over the Tor network [8] have accidentally revealed their IP address at some point in time.

The remainder of this paper is organized as follows. We discuss relevant background on OpenBazaar in Sect. 2, describe our collection methodology in Sect. 3, and present our results in Sect. 4. We explain how our work differs from related efforts in Sect. 5, and conclude in Sect. 6.

2 The OpenBazaar System

Imagine Alice wishes to join the OpenBazaar network and purchase an item. First, she downloads the OpenBazaar client from <https://openbazaar.org>, where she is presented with a “Get Started” page which allows her to edit her profile. Since all users on the platform can act as both buyers and vendors, she may also begin creating item listings if she were so inclined. Crucially, if Alice proceeds with the default setup process *even once*, her public IP address will be leaked to the network. Indeed, using Tor with OpenBazaar requires an extra configuration step, which is not immediately evident to first-time users. While the OpenBazaar client prompts the user if it detects that Tor is already running on the user’s machine, we do not expect this to be common for most users.

Once Alice creates her account, she can explore the listings on the platform. OpenBazaar’s peer-to-peer backbone is the InterPlanetary File System (IPFS),

an open-source protocol which allows content storage and addressing across a network of distributed nodes [3]. IPFS can be used as an alternative to a traditional client/server architecture – instead of clients requesting data from a server, here, a node provides a hash of the content it is requesting to its peers, who are either able to provide the requested content or query their own peers for its location.

Since IPFS does not support a centralized repository of item listings, search is implemented through a third-party search engine. At the time of writing, the only search engine enabled by default in the OpenBazaar GUI is created by the OB1 developers – other known search engines, such as SearchBizarre and a service operated by BlockStamp, require a manual addition to the client which may not be intuitive for new users. These search engines operate their own crawlers which travel the network to index users and listings, where their results can then be queried by users who have added their search engine to the OpenBazaar client. Existing third-party search engines exist mainly to provide an uncensored view of listings on the platform, as OB1’s search engine does not include listings for illicit products.

Listings are stored over IPFS using a Distributed Hash Table (DHT), where nodes store the hash of a given user, item, or feedback along with their location on the network. If a node visits a vendor page or views one of its listings, it also stores the information it receives locally for a certain period of time, allowing for an added level of redundancy across the network. This enables vendors not to be perpetually logged in. The system has been designed with the intention that individual listings will be re-seeded for about a week after their owner was last seen online [22].

When Alice clicks on a listing she is interested in, her client attempts to fetch the relevant information from her peers by querying the DHT for the item hash. If Bob, the store owner, is online, Alice will often be able to receive the listing directly from Bob – otherwise, Bob’s data is typically seeded by other IPFS peers for a set period of time. If Alice can retrieve the item hash, she is brought to a listing page which contains information about the item such as its price, description, shipping details, preferred cryptocurrency payment method, and pictures of the item. Currently, OpenBazaar supports payment in Bitcoin, Bitcoin Cash, Litecoin, and Zcash [2]. Zcash required a full installation of its binary for more than half of the duration of our study, which raises usability concerns; Litecoin was added only in the second half of our measurement interval.

When purchasing an item through OpenBazaar, Alice may opt to directly send cryptocurrency to Bob to pay him, or to use a moderator service to hold the funds in escrow until the item is received. Moderators are OpenBazaar users who volunteer to mediate disputes between other users and decide the eventual distribution of funds using multisignature transactions, and usually do so in exchange for a small fixed-percentage fee. Moderators receive community feedback, but individual moderators are chosen by the vendor at the time of listing creation. Regardless, Alice sends her payment details and shipping information to Bob

through the platform and the sale proceeds as it would on existing services such as eBay.

3 Data Collection Methodology

We turn now to our data collection methodology, outlining our objectives first, before discussing the mechanisms we use and their limitations.

3.1 Objectives

Contrary to traditional dark-web marketplaces [30], data collection on OpenBazaar is encouraged rather than discouraged, which eschews most of the constraints one faces when attempting to scrape a network stealthily. In fact, since OpenBazaar is not regulated by any central authority, our node cannot be easily banned from the network. While it could be blocked by individual nodes, we received no indication that this happened at any point during the study. Therefore, our primary focuses were on data completeness and collection speed.

We elected to scrape data ourselves, despite the existence of several search providers for OpenBazaar. Some OpenBazaar search providers service the OpenBazaar GUI (i.e., they return JSON parseable by the OpenBazaar client GUI), while others simply serve their results on public-facing webpages. Regardless of the method used, we noticed some search engines – in particular, the most popular search engine, run by OB1 at <http://openbazaar.com> – do not display results for illicit products. As a result, we needed to build our own OpenBazaar crawling infrastructure to obtain uncensored data.

Ethics of Data Collection. The data we collect are volunteered by OpenBazaar participants. In particular, listings, descriptions, and user feedback are all made publicly available for everyone to see. We previously referred to our Institutional Review Board (IRB) to determine whether data collected in related work [6, 30] could be publicly reshared. Our IRB had opined this line of work was not human-subject research. The only difference with previous research is that here, we do collect IP addresses of other clients. But, in doing so, our crawler does not collect any information a regular OpenBazaar node would not collect for operational purposes. In short, like most peer-to-peer network measurement research (see e.g., [7]), our measurements do not put users at additional risk compared to participating in the peer-to-peer network in the first place.

3.2 Crawler Design

Our crawler runs over Tor and leverages the OpenBazaar API [24], which is a modified version of the IPFS API. First, our crawler queries a list of its connected peers (GET <http://localhost:4002/ob/peers/>). For each of those peers, we retrieve their closest connected peers (GET [http://localhost:4002/ob/closestpeers/\[peer_ID\]](http://localhost:4002/ob/closestpeers/[peer_ID])), and add them to a

list, recursively continuing until we no longer find any new peers (in our testing, this usually took between five and seven rounds). Using each peer’s unique user ID, we can then make separate API calls to retrieve all of their current listings (GET `http://localhost:4002/ob/listings/[peer_ID]`) as well as all reviews left for those listings. Users, items, and reviews are all uniquely identified within OpenBazaar by a 46-character alphanumeric hash, which allows us to reliably track them across different scrapes. We log the approximate geographic location of peers with public IP addresses with FreeGeoIP (now called ipstack, [1]).

The crawler makes use of Python’s Requests library [26], and the relatively small size of the network means that it is very fast (often completing in an hour or less). However, to avoid too much redundancy and to reduce strain on the Tor network, we chose to scrape the network once every two to four hours. This provided sufficient coverage on a daily basis, and by leaving our OpenBazaar node running while the scraper was not in use, we also contributed to the overall health of the network.

3.3 Potential Limitations

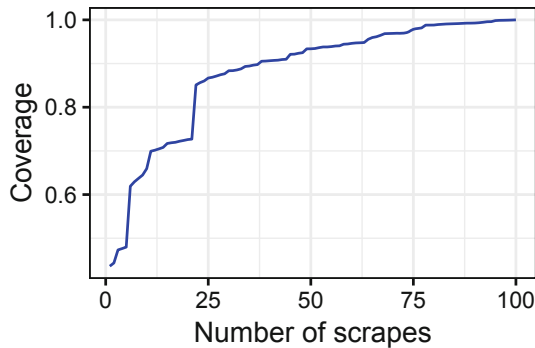


Fig. 1. Estimated coverage as a function of the number of scrapes.

Network Coverage. One risk of measuring a decentralized marketplace is that our node could lack a complete view of the network, and it was not uncommon for our scraper’s requests to the DHT to time out for certain items or users. Figure 1 depicts the relative coverage percentage of items seen on the network during April 2019. We observed 6,292 items on the network during this time, and saw 85% of them during the first two days (i.e., the first 24 scrapes), encountering the rest over the following seven. The large number of scrapes required to observe this proportion of the network is indicative of the high levels of churn we discuss in Sect. 4.1. Non-seller nodes (which make up the majority of network participants) also have little incentive to leave their OpenBazaar clients open for

long periods of time, as they do not need their listings to be seeded by other network participants. The diminishing returns shown by the curve suggest that there are few items remaining to be found by the crawler.

While the number of items we observe is less than the number reported by the OB1 developers [5], we often found through manual inspection that many items returned in the later pages of a given search result performed through their service failed to load in our client. This is likely because the OpenBazaar protocol attempts to cache data from inactive nodes for approximately a week [22] before they are “forgotten” by the network. For example, if Alice owns a store and do not log on for a few days, other nodes should still be able to access Alice’s store for approximately a week because it will be cached collectively by the other nodes in the network. Since OB1’s search engine returns results indexed by its crawlers and then stored on its own servers, it is likely that its results contain stale listings that are no longer reachable on the network.

We therefore believe that our view of the network is typical of that seen by an average node; this motivates the need for our crawler to frequently scrape the entire network to increase coverage.

Economic Coverage. Unlike many other marketplaces, leaving a review is not mandatory after a purchase on OpenBazaar. We were unable to purchase any items ourselves as a part of this study due to our legal counsel’s concerns about inadvertently participating in money laundering, but we confirmed this with one of the OpenBazaar developers. As a result, our sales numbers in Sect. 4.2 are a lower bound, even if we assume complete coverage of the network.

Despite this caveat, we are confident that our results are useful for two reasons. First, social norms on anonymous marketplaces have proven quite strong over the years: even when leaving feedback is not mandatory, buyers often do so, especially if they plan on buying from the same seller again in the future [30]. Soska and Christin’s analysis, based on feedback ratings [30], showed numbers very close to those obtained externally through criminal complaints, when vendors were arrested or through marketplace takedowns, even when leaving feedback was not actively enforced by all marketplaces. The importance of feedback on those marketplaces was also evident when the original Silk Road changed the way feedback was tallied (shifting from per-item feedback, to aggregate, per-vendor feedback) and quickly reversed course in the face of customer complaints [6].

Furthermore, our reported distribution of sales by category is likely valid. One could think that users purchasing illicit items on OpenBazaar would be less likely to leave a public review than users purchasing legal products. As we will see, our results do not support this hypothesis.

4 Results

We next turn to our measurement results. We first describe network-level metrics such as the size of the peer-to-peer network, or the underlying churn in IP

addresses we see participating in the network, before turning to discussing the economic activity that appears to take place on OpenBazaar. We then examine the security and privacy precautions vendors take.

4.1 Network-Level Metrics

The OpenBazaar Peer-to-Peer Network. Over our entire measurement interval, we have observed 6,651 distinct network participants, using 6,116 distinct IP addresses.

Table 1. OpenBazaar demographics. Users do not have any product listings; sellers have at least one active listing; active sellers have realized at least one documented sale.

	With Tor	Without Tor	Total
Users	743	4,487	5,230
Sellers	197	1,057	1,254
Active sellers	40	127	167
Total	980	5,670	6,651

Table 1 breaks down these participants into finer-grained demographics, distinguishing between *users*, who do not list any product, and are therefore assumed to be solely browsing or buying items; *sellers*, who list at least one product, but do not have any documented sale – that is no one left any feedback for them; and *active sellers*, who have received at least one piece of feedback. We also break down these participants between those who use Tor and those who do not.

Figure 2 shows, for each day, the number of OpenBazaar hosts we have encountered during our multiple scrapes on that day. The lighter curve denotes the fraction of hosts that are using the Tor network. We observe that the population has been relatively steady, at approximately 80 users simultaneously online on any given day throughout our measurement interval. The few downward spikes denote measurement issues rather than network instability.

The gaps in the plot denote times during which our measurement infrastructure was disabled, or otherwise unable to properly collect data.

The seemingly decreasing number of Tor users toward the end of the measurement interval might be a slight undercount. Through experimentation with our own test nodes, we discovered that OpenBazaar nodes in version 2.3.1 and higher running over Tor, sometimes failed to appear in our crawls, despite being online and reachable (that is, if one knew their node ID). This coincides, and may be due to the backwards-incompatible [23] software upgrade on March 19, 2019 (OpenBazaar 2.3.1), which may have affected some long-term participants over Tor. Overall, compared to Table 1, this plot seems to indicate that the vast majority of OpenBazaar users are actually rarely online. To better understand

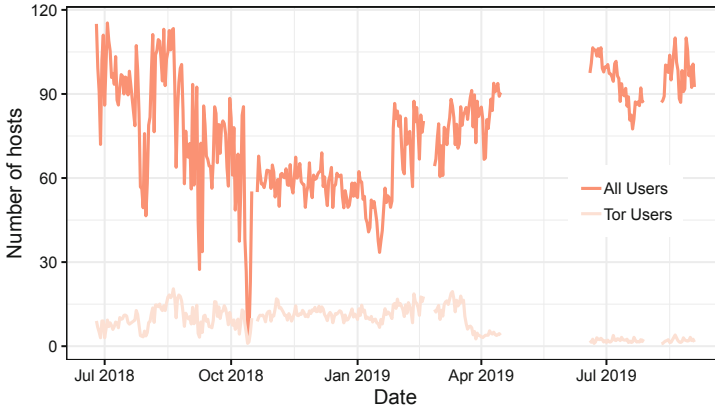


Fig. 2. Number of hosts on the OpenBazaar peer-to-peer network. Each point is the cumulative number of different hosts seen over all measurements taken on a specific day.

the dynamics of the OpenBazaar network, we next turn to a survivability analysis. Similar to Christin [6], we estimate vendor “lifespan” by recording the time we first saw a vendor, and the time we saw them last. They may have left and rejoined in the meantime—here we are looking at the vendor lifetime, regardless of their transient activity. To account for measurement effects (e.g., vendors still being present on the last day of measurement), Fig. 3 depicts a Kaplan-Meier estimator [16] that shows the probability a given user seen on day 0 will be seen again after x days, broken down by the categories defined in Table 1.

Churn is high among regular users: more than two thirds of them stay less than a day. Vendors, on the other hand tend to stay longer, especially vendors that have documented sales (i.e., that have received feedback). Roughly three quarters of all participants do not stay more than a week; a few users remain on the network throughout our measurement interval. Log-rank tests confirm that visually striking differences between the survival curves for all of these user categories are statistically significant ($p < 0.0001$). In short: vendors tend to be long-lived, while regular users are not, and usage of Tor is positively correlated with longer presence on the network.

Geographical Considerations. Figure 4 shows the geolocation of the IP addresses of participants that are not using Tor. OpenBazaar seems to be fairly international, with the usually observed concentration of users in North America and Europe. A couple of points with strange locations (e.g., Easter Island) suggest certain participants use VPNs, some of which are known to advertise implausible locations that fool geolocation databases [32]. We generally do not observe meaningful differences between different types of users, even though Western Africa,

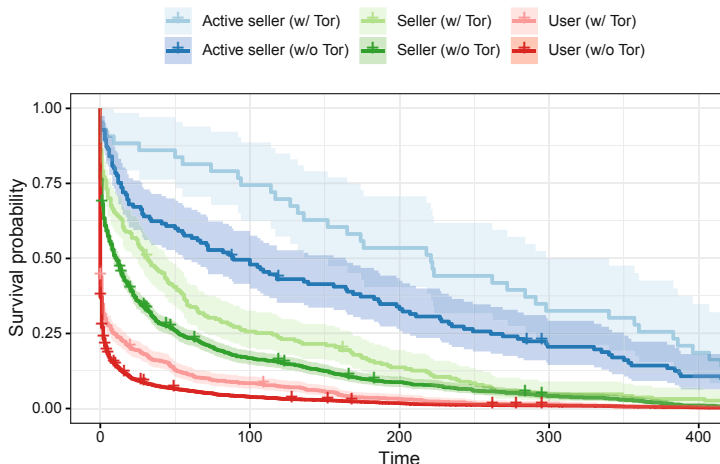


Fig. 3. Survivability analysis of OpenBazaar users. Kaplan-Meier estimator that shows the probability a given user seen on day 0 will be seen again after x days, broken down for different types of users. Shaded areas indicate 95% C.I. (Color figure online)

India and Thailand/Malaysia feature a larger proportion than we expected of active sellers – again, it is hard to tell whether this could be due to VPN usage.

4.2 Economic Activity

We next turn to a study of economic activity on OpenBazaar. We observed 24,379 distinct item listings during our measurement intervals. The apparent high average ratio of listings per seller (≈ 40) is due to the ease of creating listings (including test listings) and to a few “power users” who have thousands of listed items on the platform.

Item Listing Survivability. We start with a survival plot in Fig. 5 shows the probability an item seen at time zero will still be available on the network x days later. The median item stays online for approximately three weeks. One quarter of all items disappears (i.e., are delisted) within a day, which further motivates the need to repeatedly crawl the network for completeness. A handful of items were present throughout our measurement interval. The “jumps” observed in the graph correspond to a large number of items belonging to a given vendor being all delisted on the same day, presumably because the vendor node had been offline for long enough to have its listings cleared from the IPFS cache.¹

¹ Item listings do *not* automatically disappear when an item is sold out. The seller needs to either delete the listing, or be offline for a long enough period of time for the listing to stop being seeded by the network. We differentiate between the two cases by only counting visible reviews for an item as a sale and not considering disappearing listings as possible sales.

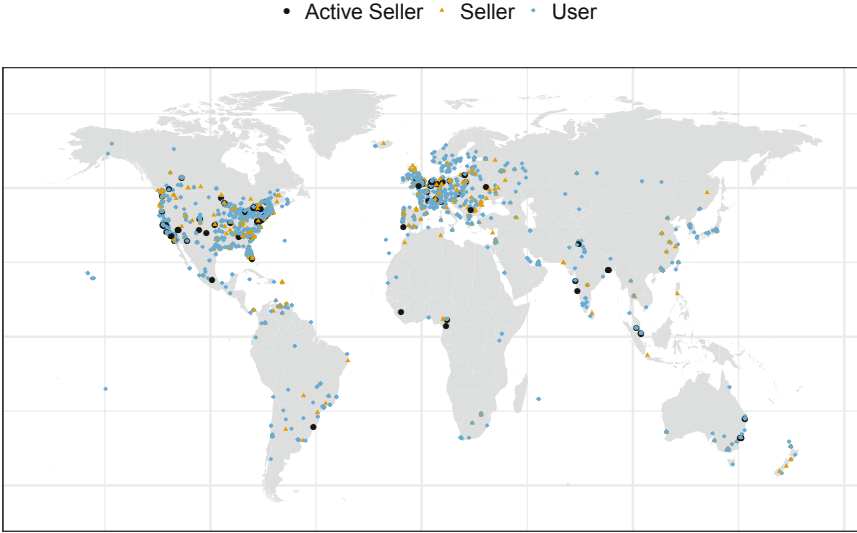


Fig. 4. Geolocation of OpenBazaar users. Users are network participants; Sellers are users with at least one listing; Active sellers are sellers with at least one feedback.

Overall, the survivability analysis paints a picture eerily similar to that of the early days of online anonymous marketplaces [6]: most items are very short-lived. As discussed in prior work [6,30], short-lived items are usually indicative of vendors holding low stocks, which in turn suggests that vendors operate primarily in the retail space, with small product quantities, and (usually) low sales volumes.

Preliminary Economic Analysis. We next examine economic activity on OpenBazaar, across item categories. We initially attempted to feed each of the 24,379 item listings we observed into the 16-category item classifier proposed by Soska and Christin [30], who trained their item classification with listings from the Agora and Evolution marketplaces and showed very high (>98%) accuracy when evaluating *a priori* unknown listings.

We realized, however, that items classified in the “Miscellaneous” category represented a disproportionately high fraction of all listings, which led us to further break this category down into additional categories: Adult Toys, Art, Clothing, Jewelry, Print Media, and Souvenirs. We took 200 new items from our OpenBazaar corpus, hand-labeled them,² and added them five times (i.e., an extra 1,000 items) to the original training set consisting of 62,989 labeled items from Evolution and Agora. Our resulting classifier operates on 22 categories. Table 2 shows very good performance metrics, with precision and recall,

² A single researcher was tasked with this labeling, hence we do not report agreement numbers.

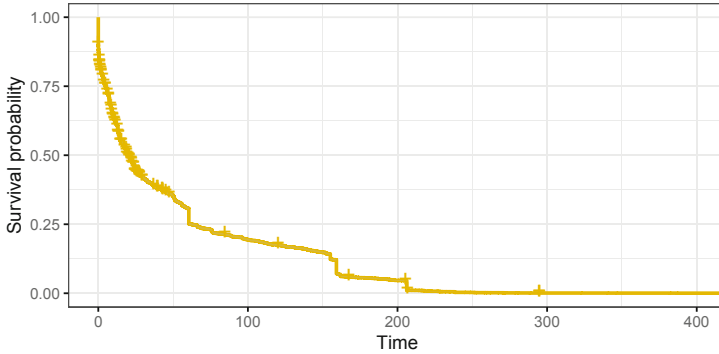


Fig. 5. Survivability analysis of OpenBazaar items. The plot is a Kaplan-Meier estimator that shows the probability a given item seen on day 0 will be seen again after x days. The median lifetime of an item is approximately 22 days. The curve fully overlaps the very narrow 95% confidence interval.

overall, being over 0.96. This is unsurprising, given that the modified classifier is very close to the original classifier from Soska and Christin – we merely added $1,000/63,989 \approx 1.6\%$ of new training data to the original classifier. The support, here, is much larger than the number of items we observed in OpenBazaar, since we are evaluating with 10-fold cross-validation on the original dataset provided by Soska and Christin³ [30] and the OpenBazaar data. Certain categories (e.g., Adult Toys) do appear very rarely, though. The overall support for the new categories is fairly small (240 items); this again is unsurprising, as the imbalance in training sets means that an item needs to closely resemble a training example to be classified as one of the new sub-categories. In other words, our modified classifier closely mimics what Soska and Christin used; in a few “obvious” cases, it will manage to further identify a subcategory, but will do so very conservatively.

Next, we present the category breakdown for these 24,379 listings in Table 3. Simply looking at listing counts (columns 2 and 3), close to half of the listings are in the “Misc.,” “Print Media,” and “Souvenirs” categories, which generally denote legitimate items. The third largest category, “Digital goods,” contains a mix of legitimate (e.g., e-books and other guides) and illegitimate (pornographic website account passwords) items. In short, a majority of items *listed* on OpenBazaar shops appear to be for legal products.

However, looking at actual sales (columns 4 through 6) paints a very different picture. To compute sales for a given item listing, using the same technique as in related efforts [6, 30], we add up the price of the item to its total sales at each time a feedback for a sale of that item is recorded. Feedback that were recorded prior to our monitoring the OpenBazaar network are counted if the corresponding item was still listed when we scraped the network, and feedback

³ The dataset is available from IMPACT, <https://www.impactcybertrust.org>.

Table 2. Classifier performance.

Category	Precision	Recall	F1-score	Support
Adult Toys	1.00	1.00	1.00	3
Art	1.00	1.00	1.00	24
Benzos	0.97	0.98	0.97	21,132
Cannabis	1.00	1.00	1.00	113,516
Clothing	1.00	1.00	1.00	38
Digital Goods	0.94	0.96	0.95	158,162
Dissociatives	1.00	0.99	0.99	7,172
Drug Paraphernalia	0.97	0.98	0.97	15,740
Ecstasy	1.00	0.99	0.99	49,184
Electronics	0.96	0.94	0.95	5,379
Jewelry	0.76	1.00	0.87	13
Misc	0.87	0.80	0.83	47,651
Opioids	0.98	0.98	0.98	25,511
Prescription	0.95	0.93	0.94	23,023
Print Media	1.00	0.96	0.98	113
Psychedelics	1.00	1.00	1.00	31,023
Sildenafil	0.98	0.97	0.97	31,22
Souvenirs	1.00	0.90	0.95	49
Steroids	0.99	1.00	0.99	17,044
Stimulants	0.98	0.99	0.99	55,555
Tobacco	1.00	1.00	1.00	3,966
Weapons	0.99	0.97	0.98	5,341
Total	0.96	0.96	0.96	582,761

was still accessible. As noted before, reviews are not mandatory in OpenBazaar, so that the sales numbers presented are a lower bound. Even with this caveat, sales volumes appear to be very modest. We only count around \$217,000 in total sales over our measurement interval. By comparison, sales on Silk Road [6, 30] and AlphaBay [21], two of the largest online anonymous marketplaces, reached hundreds of *millions* of dollars in revenue.

Looking a bit more carefully at the data revealed an interesting outlier: one vendor, *P...A*, appeared to be single-handedly responsible for 60% of all sales on OpenBazaar. In particular, that vendor had three pieces of feedback for a \$36,000+ listing for a kilogram of cocaine, which should have accounted for more than \$100,000 in sales by itself. When we manually inspected this specific vendor, we discovered that for most of their items, the feedback appeared to be fake: very closely clustered timestamps, all with highly positive ratings and uninformative messages for highly priced items. This strongly suggested an attempt at manip-

Table 3. Sales observed during our measurement interval. The values in columns 4–6 exclude vendor *P...A*, whose feedback seems highly questionable, and likely fake.

Category	Listings (count)	Listings (%)	Sales (count, corrected)	Sales (USD, corrected)	Sales (%, corrected)
Adult Toys	182	0.747	0	0	0.0
Art	331	1.358	10	975	1.13
Benzos	155	0.636	8	1,650	1.92
Cannabis	1,910	7.835	318	22,450	26.1
Clothing	1,881	7.72	30	375	0.437
Digital Goods	2,701	11.079	139	1,312	1.53
Dissociatives	21	0.086	2	1,050	1.22
Drug Paraphernalia	829	3.4	13	224	0.261
Ecstasy	243	0.997	10	4,895	5.69
Electronics	1,906	7.818	23	2,389	2.78
Jewelry	354	1.452	9	208	0.243
Misc	6,796	27.876	333	4,028	4.69
Opioids	223	0.915	17	2,207	2.57
Prescription	289	1.185	23	808	0.941
Print Media	4,902	20.107	6	123	0.144
Psychedelics	242	0.993	125	17,653	20.5
Sildenafil	45	0.185	41	792	0.921
Souvenirs	791	3.245	60	3,587	4.17
Steroids	69	0.283	2	11	0.0128
Stimulants	242	0.993	24	20,835	24.2
Tobacco	34	0.139	1	5	0.005
Weapons	233	0.956	8	374	0.436
Total	24,379	100	1,202	85,954	100

ulation by the vendor.⁴ While in centralized marketplaces, padding feedback with misleading information is prohibited, and frequently results in banning the vendors engaging in such deceptive practices, the decentralized nature of OpenBazaar makes this kind of enforcement difficult. We do note, though, that OpenBazaar supports moderators that can assist in ensuring transactions are conducted satisfactorily (see Sect. 2); unsurprisingly, all of *P...A*'s listings were unmoderated. We removed this vendor, and the 33 associated sales, from further consideration in columns 4–6, to obtain an hopefully more accurate picture of

⁴ Interestingly, *one* of their items seemed to have legitimate feedback, which pre-dated all of the seemingly deceptive feedback discussed here.

sales on OpenBazaar – excluding *P...A*, the total amount of sales is actually around \$86,000.

As Table 3 shows, the category distribution of items that do sell is very different from the category distribution of items that are merely listed. Over 25% of all recorded sales are for cannabis-related products (including seeds), and more than three quarters of all recorded sales are for drugs – prescription drugs or narcotics.⁵ This higher economic activity occurs despite the fact that the OpenBazaar developers have taken active measures to try to make illicit items harder to find, by excluding them from their built-in search engine search results, which suggests that the demand for illicit offerings far outpaces that for legitimate goods available on OpenBazaar.

Table 4. Distribution of feedback ratings. 5 is best, 1 is worst.

Score	Count	Percent
5	1,302	85.32%
4	28	1.83%
3	34	2.22%
2	29	1.90%
1	133	8.71%

Feedback Ratings. An alternative hypothesis would be that buyers of legitimate products are somehow less likely to leave feedback than buyers of illicit goods. We found no evidence to support that hypothesis. Specifically, we present the feedback ratings we observed in Table 4. OpenBazaar uses a 5-point rating scale, where higher scores are better, i.e., vendor strive to obtain 5-star feedback. The ratings we see heavily skew positive, as has been observed in general (legitimate) e-commerce platforms [27], and feedback distribution presents striking similarities with that obtained for feedback left by Silk Road patrons [6, Table 3]: 5's dominate, followed by 1's, and other ratings are less frequently used.

4.3 Operational Security

While much of the core OpenBazaar vendor base tends to connect over Tor, not all of these users are truly anonymous. Indeed, user IDs are persistent across sessions. By comparing the unique user IDs of the 980 nodes seen connecting over Tor with those seen connecting over public IP addresses, we found that 173 users (17.7%) had revealed an IP address at some point in time. Not all of these users may have wished to remain anonymous during the entire collection cycle, but we did observe some obvious lapses in operational security, such as US vendors selling marijuana offering global shipping. It appears a version of

⁵ The disproportionate volume of Psychedelics sales seems to be mostly influenced by one specific vendor, who has been highly successful on various online anonymous markets, and also operates their own vendor shop, and has presence on OpenBazaar.

the OpenBazaar client which is pre-bundled with Tor is in development, which should alleviate this problem over time [25].

At the beginning of our measurement timeframe, OpenBazaar nodes were tied to a single cryptocurrency – if a vendor wished to, for example, sell the same items both with Bitcoin and Bitcoin Cash, they were required to maintain two distinct nodes with identical listings. Furthermore, for some time using Zcash on the platform also required running a full Zcash node. As a result, activity on the platform was conducted almost entirely in Bitcoin until the release of a multiwallet feature on January 17, 2019, which added native support for the three currencies mentioned plus Litecoin. Previous studies (e.g., [20]) have shown that Bitcoin is highly traceable, meaning that early purchases on the platform may be able to de-anonymize certain vendors. Since the introduction of the multiwallet, however, there has been large growth in Zcash usage, with at least⁶ 12,441 observed listings accepting payment in the currency.

5 Related Work

This work follows a long lineage of peer-to-peer network measurements, which can be traced back to studies of file-sharing networks such as Napster [29], or Gnutella [28]; or, beyond file-sharing, of Skype [11]. Later papers focused on specific metrics to better understand user behavior. For instance, Gummadi et al. studied peer churn on Kazaa [12], while others investigated overall peer availability [4, 31], or resilience to poisoning attacks [7, 19]. A number of papers looked into the economics of online anonymous marketplaces and have documented their rise in popularity [6, 9, 30]. In comparison, our analysis of OpenBazaar shows fairly modest revenues. Finally, also related to the operational security aspects we outline in this paper are attempts to quantify cryptocurrency traceability – notably efforts to trace Bitcoin [20], Monero [21], and Zcash [17].

6 Conclusion

We conducted multiple daily crawls of the OpenBazaar distributed marketplace over approximately 14 months (June 25, 2018–September 3, 2019). More than half of the 6,651 participants we observed were present only for less than a day, but users relying on Tor tended to be much longer lived, particularly if they were actively selling items. Economic activity is orders of magnitude smaller than on centralized anonymous marketplaces, and while most listed items are for legitimate products, the majority of items that do result in sales are for narcotics. Finally, vendors are not always using prudent security practices, leaking for instance their IP address despite generally connecting over Tor, which makes a strong case for imposing secure defaults—fortunately, the OpenBazaar developers are already reportedly working on this [25].

⁶ A bug in our parser code for items with multiple currencies prevented us from precisely computing the number of such listings, but we could recover this lower bound.

We conclude that OpenBazaar, so far, has not gained much traction to usher in the new era of decentralized, private, and legitimate electronic commerce it was promising.⁷ This could be due to a lack of user demand for decentralized marketplaces, lack of integration of private features, or other factors, such as a higher learning curve for users compared to centralized alternatives.

Acknowledgments. We thank our shepherd, Ben Edwards, and the anonymous reviewers for numerous comments that greatly improved this paper. We are also grateful to Kyle Soska for extensive discussions about this work. This research was partially supported by DHS Office of Science and Technology under agreement number FA8750-17-2-0188.

References

1. Apilayer: ipstack (2019). <https://ipstack.com/>. Accessed 20 June 2020
2. Ben-Sasson, E., et al.: Zerocash: Decentralized anonymous payments from bitcoin. In: IEEE Symposium on Security and Privacy, pp. 459–474. IEEE Computer Society, May 2014
3. Benet, J.: IPFS - content addressed, versioned, P2P file system. Technical report, June 2014, draft 3. <https://arxiv.org/abs/1407.3561>
4. Bhagwan, R., Savage, S., Voelker, G.M.: Understanding availability. In: Kaashoek, M.F., Stoica, I. (eds.) IPTPS 2003. LNCS, vol. 2735, pp. 256–267. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45172-3_24
5. CC_EF_JTF: Did OpenBazaar ever take off? OpenBazaar developer comments on total number of listings, April 2019. https://www.reddit.com/r/Bitcoin/comments/bad0nh/did_openbazaar_ever_take_off/ekbd2k6/
6. Christin, N.: Traveling the silk road: a measurement analysis of a large anonymous online marketplace. In: Proceedings of the 22nd World Wide Web Conference (WWW 2013), Rio de Janeiro, Brazil, pp. 213–224, May 2013
7. Christin, N., Weigend, A., Chuang, J.: Content availability, pollution and poisoning in peer-to-peer file sharing networks. In: Proceedings of ACM EC 2005, Vancouver, BC, Canada, pp. 68–77, June 2005
8. Dingledine, R., Mathewson, N., Syverson, P.: Tor: the second-generation onion router. In: Proceedings of the 13th USENIX Security Symposium, August 2004
9. European Monitoring Centre for Drugs and Drug Addiction and Europol: Drugs and the darknet: perspectives for enforcement, research and policy, November 2017. <http://www.emcdda.europa.eu/system/files/publications/6585/TD0417834ENN.pdf>
10. Greenberg, A.: Inside the ‘darkmarket’ prototype, a Silk Road the FBI can never seize, April 2014. <https://www.wired.com/2014/04/darkmarket/>
11. Guha, S., Daswani, N.: An experimental study of the skype peer-to-peer VoIP system. Technical report, Cornell University (2005)
12. Gummadi, K., Dunn, R., Saroiu, S., Gribble, S., Levy, H., Zahorjan, J.: Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In: Proceedings of ACM SOSP 2003, Bolton Landing, NY, pp. 314–329, October 2003

⁷ We sent an earlier preprint of this paper to the OpenBazaar developers, but did not receive any feedback.

13. Higgins, S.: OpenBazaar raises \$1 million for decentralised marketplace, June 2015. <https://www.coindesk.com/openbazaar-raises-1-million-from-silicon-valley-giants>
14. Higgins, S.: Bitcoin-powered marketplace OpenBazaar raises \$3 million, December 2016. <https://www.coindesk.com/bitcoin-powered-marketplace-openbazaar-raises-3-million>
15. Jeffries, J.: Remove BlockBooth from search providers. OpenBazaar developer comments on OB1 censorship policy, April 2019. <https://github.com/OpenBazaar/openbazaar-desktop/issues/1569>
16. Kaplan, E., Meier, P.: Nonparametric estimation from incomplete observations. *J. Am. Stat. Assoc.* **53**, 457–481 (1958)
17. Kappos, G., Yousaf, H., Maller, M., Meiklejohn, S.: An empirical analysis of anonymity in Zcash. In: Proceedings of the USENIX Security, August 2018
18. Kruithof, K., Aldridge, J., Décarry Hétu, D., Sim, M., Dujso, E., Hoorens, S.: Internet-facilitated drugs trade: an analysis of the size, scope and the role of the Netherlands. Technical report RR1607 (2016). https://www.rand.org/pubs/research_reports/RR1607.html
19. Liang, J., Kumar, R., Xi, Y., Ross, K.: Pollution in P2P file sharing systems. In: Proceedings of IEEE INFOCOM 2005, Miami, FL, March 2005
20. Meiklejohn, S., et al.: A fistful of bitcoins: characterizing payments among men with no names. In: Proceedings of the ACM/USENIX Internet Measurement Conference, Barcelona, Spain, pp. 127–140, October 2013
21. Möser, M., et al.: An empirical analysis of traceability in the Monero blockchain. In: Proceedings of PETS, Barcelona, Spain, vol. 3, July 2018
22. OB1 Team: OpenBazaar seller guide - what to expect in this decentralized marketplace, October 2017. <https://openbazaar.org/blog/openbazaar-seller-guide-what-to-expect-in-this-decentralized-marketplace/>
23. OpenBazaar: OpenBazaar 2.3.1 is released with IPFS rebase, March 2019. <https://www.openbazaar.org/blog/openbazaar-2-3-1-release-ipfs-rebase/>
24. OpenBazaar Team: Openbazaar API. <https://api.docs.openbazaar.org/>
25. Patt, S.: Create “OpenBazaar-Private” installer, March 2019. <https://github.com/OpenBazaar/openbazaar-desktop/issues/1735>
26. Reitz, K.: Requests: Http for humans, May 2019. <https://requests.readthedocs.io>
27. Resnick, P., Zeckhauser, R.: Trust among strangers in internet transactions: empirical analysis of eBay’s reputation system. *Advances in Applied Microeconomics*, vol. 11, pp. 127–157 (2002)
28. Ripeanu, M., Foster, I.: Mapping the gnutella network: macroscopic properties of large-scale peer-to-peer systems. In: Druschel, P., Kaashoek, F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 85–93. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45748-8_8
29. Saroiu, S., Gummadi, K., Gribble, S.: A measurement study of peer-to-peer file sharing systems. In: Proceedings of SPIE/ACM MMCN 2002, San Jose, CA, pp. 156–170, January 2002
30. Soska, K., Christin, N.: Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In: Proceedings of the 24th USENIX Security Symposium (USENIX Security 2015), Washington, DC, pp. 33–48, August 2015
31. Stutzbach, D., Rejaie, R.: Understanding churn in peer-to-peer networks. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, pp. 189–202. ACM (2006)
32. Weinberg, Z., Cho, S., Christin, N., Sekar, V., Gill, P.: How to catch when proxies lie: verifying the physical locations of network proxies with active geolocation. In: Proceedings of the 17th ACM Internet Measurement Conference (IMC 2017), Boston, MA, October 2018