

Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace

Nicolas Christin

July 30, 2012

(Revised November 28, 2012)

[CMU-CyLab-12-018](#)

[CyLab](#)

Carnegie Mellon University
Pittsburgh, PA 15213

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 28 NOV 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012			
4. TITLE AND SUBTITLE Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University, CyLab, Pittsburgh, PA, 15213		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We perform a comprehensive measurement analysis of Silk Road, an anonymous, international online marketplace that operates as a Tor hidden service and uses Bitcoin as its exchange currency. We gather and analyze data over eight months between the end of 2011 and 2012, including daily crawls of the marketplace for nearly six months in 2012. We obtain a detailed picture of the type of goods being sold on Silk Road, and of the revenues made both by sellers and Silk Road operators. Through examining over 24,400 separate items sold on the site, we show that Silk Road is overwhelmingly used as a market for controlled substances and narcotics, and that most items sold are available for less than three weeks. The majority of sellers disappears within roughly three months of their arrival, but a core of 112 sellers has been present throughout our measurement interval. We evaluate the total revenue made by all sellers, from public listings, to slightly over USD 1.2 million per month; this corresponds to about USD 92,000 per month in commissions for the Silk Road operators. We further show that the marketplace has been operating steadily, with daily sales and number of sellers overall increasing over our measurement interval. We discuss economic and policy implications of our analysis and results including ethical considerations for future research in this area.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	27	

Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace

Nicolas Christin
Carnegie Mellon INI/CyLab
nicolasc@cmu.edu

Working paper

First version: May 4, 2012.

This version: November 28, 2012.

Id: paper.tex 1654 2012-11-28 17:04:37Z nicolasc

Abstract

We perform a comprehensive measurement analysis of Silk Road, an anonymous, international online marketplace that operates as a Tor hidden service and uses Bitcoin as its exchange currency. We gather and analyze data over eight months between the end of 2011 and 2012, including daily crawls of the marketplace for nearly six months in 2012. We obtain a detailed picture of the type of goods being sold on Silk Road, and of the revenues made both by sellers and Silk Road operators. Through examining over 24,400 separate items sold on the site, we show that Silk Road is overwhelmingly used as a market for controlled substances and narcotics, and that most items sold are available for less than three weeks. The majority of sellers disappears within roughly three months of their arrival, but a core of 112 sellers has been present throughout our measurement interval. We evaluate the total revenue made by all sellers, from public listings, to slightly over USD 1.2 million per month; this corresponds to about USD 92,000 per month in commissions for the Silk Road operators. We further show that the marketplace has been operating steadily, with daily sales and number of sellers overall increasing over our measurement interval. We discuss economic and policy implications of our analysis and results, including ethical considerations for future research in this area.

Keywords: Online crime, anonymity, electronic commerce.

Welcom
messages(0) | orders(0) | account(\$0.00) | settings |

1 day 00 hrs 00 mins 00 secs until Four Twenty!!!

Shop by category:
 Drugs(2788)
 Cannabis(796)
 Dissociatives(48)
 Ecstasy(307)
 Opioids(211)
 Other(98)
 Prescription(541)
 Psychedelics(366)
 Stimulants(235)
 Apparel(28)
 Books(286)
 Computer equipment(13)
 Digital goods(219)
 Drug paraphernalia(74)
 Electronics(17)
 Fireworks(1)

170\$ pecunix \$39.23	1 OZ of Jamaican Oil \$73.91	Need bitcoins? Bitcoins for your... \$0.00
20 Grams of MDMA crystals \$124.60	HYDRO 10/325 NORCO/LORATAB \$1.75 ...	1oz - "Swazi Red" (Rooibaard)... \$29.61

News:

- Who's your favorite?
- Acknowledging Heroes
- A new anonym market **The Armory!**
- State of the R Address

Figure 1: **Silk Road front page.** The site offers a number of licit and illicit items, with a marked focus on narcotics.

1 Introduction

“More brazen than anything else by light-years” is how U.S. Senator Charles Schumer characterized Silk Road [5], an online anonymous marketplace. While a bit of a hyperbole, this sentiment is characteristic of a certain nervousness among political leaders when it comes to anonymous networks. The relatively recent development of usable interfaces to anonymous networks, such as the “Tor browser bundle,” has indeed made it extremely easy for anybody to browse the Internet anonymously, regardless of their technical background. In turn, anonymous online markets have emerged, making it quite difficult for law enforcement to identify buyers and sellers. As a result, these anonymous online markets very often specialize in “black market” goods, such as pornography, weapons or narcotics.

Silk Road is one such anonymous online market. It is not the only one – others, such as Black Market Reloaded [3], the Armory [1], or the General Store [7] are or have been offering similar services – but it gained fame after an article posted on Gawker [10], which resulted in it being noticed by congressional leaders, who demanded prompt action be taken. It is also reportedly very large, with estimates mentioned in the Silk Road online forum [6] ranging between 30,000 and 150,000 active customers.

Figure 1 shows the Silk Road front page. The site has a professional, if minimalist, look, and appears to offer a variety of goods (e.g., books, digital goods, digital currency...), but seems to have a clear focus on drugs. Not only do most items listed appear to be controlled substances, but the screenshot also shows the site advertising a sale campaign for April 20 – also known as “Pot day” due to the North American slang for cannabis (four-twenty).

In this paper, we try to provide a scientific characterization of the Silk Road marketplace, by gathering a

set of controlled measurements over roughly six months (February 3, 2012 – July 24, 2012), and analyzing them.

Specifically, we offer the following contributions. We devise a (simple) collection methodology to obtain publicly available Silk Road market data. We use the data collected to characterize the items being sold on Silk Road and the seller population. We describe how items sold and seller population have evolved over time. Using (mandatory) buyer feedback reports as a proxy for sales, we characterize sales volumes over our measurement interval. We provide an estimate of the daily dollar amount of sales conducted on Silk Road, and use this estimate to infer the amount collected in commission by Silk Road operators. While we cannot estimate the number of buyers with the dataset we collect, we show that Silk Road is a relatively significant market, with a few hundred sellers, and monthly total revenue of about USD 1.2 million. We also show that Silk Road appears to be growing over time, albeit not at the exponential rate that is claimed in forums [6].

The rest of this paper is structured as follows. We start by describing how Silk Road operates in Section 2. We then explain how we gather our measurements in Section 3. We report on our measurements analysis in Section 4, before turning to economic implications in Section 5. We discuss our findings, reflect on possible intervention policies, and ethical considerations in Section 6, outline related work in Section 7, and conclude in Section 8.

2 Silk Road overview

Silk Road is an online anonymous marketplace that started its operations in February 2011 [6]. Silk Road is not, itself, a shop. Instead, it provides infrastructure for sellers and buyers to conduct transactions in an online environment. In this respect, Silk Road is more similar to Craigslist, eBay or the Amazon Marketplace than to Amazon.com. The major difference between Silk Road and these other marketplaces is that Silk Road focuses on ensuring, as much as possible, anonymity of both sellers and buyers. In this section, we summarize the major features of Silk Road through a description of the steps involved in a typical transaction: accessing Silk Road, making a purchase, and getting the goods delivered.

Accessing Silk Road. Suppose that Bob (B), a prospective buyer, wants to access the Silk Road marketplace (SR). Bob will first need to install a Tor client on his machine, or use a web proxy to the Tor network (e.g. <http://tor2web.org>) as Silk Road runs only as a Tor hidden service [11]. That is, instead of having a DNS name mapping to a known IP address, Silk Road uses a URL based on the pseudo-top level domain `.onion`, that can only be resolved through Tor. At a high level, when Bob’s client attempts to contact the Silk Road server URL (<http://silkroadvb5piz3r.onion> at the time of this writing), Tor nodes set up a rendez-vous point inside the Tor network so that the client and server can communicate with each other while maintaining their IP addresses unknown from observers and from each other.

Once connected to the Silk Road website, Bob will need to create an account. The process is simple and merely involves registering a user name, password, withdrawal PIN, and answering a CAPTCHA. After this registration, Bob is presented with the Silk Road front page (see Figure 1) from where he can access all of Silk Road’s public listings.

Public and stealth listings. Silk Road places relatively few restrictions on the types of goods sellers can offer. From the Silk Road sellers’ guide [5],

“Do not list anything who’s (sic) purpose is to harm or defraud, such as stolen items or info, stolen credit cards, counterfeit currency, personal info, assassinations, and weapons of any kind. Do not list anything related to pedophilia.”

Conspicuously absent from the list of prohibited items are prescription drugs and narcotics, as well as adult pornography and fake identification documents (e.g., counterfeit driver’s licenses). Weapons and ammunition used to be allowed until March 4, 2012, when they were transferred to a sister site called The Armory [1], which operated with an infrastructure similar to that of Silk Road. Interestingly, the Armory closed in August 2012 reportedly due to a lack of business [6].

Not all of the Silk Road listings are public. Silk Road supports *stealth listings*, which are not linked from the rest of Silk Road, and are thus only accessible by buyers who have been given their URL. Stealth listings are frequently used for *custom listings* directed at specific customers, and established through out-of-band mechanisms (e.g., private messaging between seller and buyer). Sellers may further operate in *stealth mode*, meaning that their seller page and *all* the pages of the items they have for sale are not linked from other Silk Road pages. While Silk Road is open to anybody, stealth mode allows sellers with an established customer base to operate their business as invitation-only.

Making a purchase. After having perused the items available for sale on Silk Road, Bob decides to make a purchase from Sarah (S), a seller. While Tor ensures communication anonymity, Silk Road needs to also preserve payment anonymity. To that effect, Silk Road only supports Bitcoin (BTC, [30]) as a trading currency. Bitcoin is a peer-to-peer, distributed payment system that offers its participants to engage in verifiable transactions without the need for a central third-party. Bob thus needs to first procure Bitcoins, which he can do from the many online exchanges such as Mt.Gox [4]. Once Bob has Bitcoins, and decides to purchase the item from Sarah, instead of paying Sarah directly, Bob places the corresponding amount in *escrow* with Silk Road. Effectively, B pays SR , not S . The escrow mechanism allows the market operator to accurately compute their commission fees, and to resolve disputes between sellers and buyers. Silk Road mandates all sellers and buyers use the escrow system. Failure to do so is punishable by expulsion from the marketplace [5].

Finalizing. Once the purchase has been made, Sarah must ship it to Bob. Thus, Sarah needs a physical address where to send the item. To preserve anonymity, Silk Road recommends to use delivery addresses that are distinct from the buyer’s residence. For instance, Bob could have the item delivered at Patsy’s house, or to a post-office box. Once Sarah has marked the item as shipped, Bob’s delivery address is erased from all records. Once the item reaches its destination, Bob *finalizes* the purchase, that is, he tells Silk Road to release the funds held in escrow to Sarah (i.e., SR now pays S), and leaves feedback about Sarah. Finalizing is mandatory: if Bob forgets to do so, Silk Road will automatically finalize pending orders after a set amount of time.

Sellers with more than 35 successful transactions and who have been active for over a month are allowed to ask their buyers to *finalize early*; that is, to release payment and leave feedback before they actually receive the item. Due to the potential for abuse, Silk Road discourages finalizing early in general, and prohibits it for new sellers.

Finally, Silk Road enhances transaction anonymity by providing “tumbler” services that consist of inserting several dummy, single-use intermediaries between a payer and a payee. That is, instead of having a payment correspond to a simple transaction chain $B \rightarrow SR \rightarrow S$, the payment goes through a longer chain $B \rightarrow I_1 \rightarrow \dots \rightarrow I_n \rightarrow S$ where (I_1, \dots, I_n) are one-time-use intermediaries.

3 Collection methodology

We next turn to describing how we collected measurements of the Silk Road marketplace. We first briefly explain our crawling mechanism, before outlining some of the challenges we faced with data collection. We

then discuss in detail the data that we gathered.

3.1 Crawling mechanism

We registered an account on Silk Road in November 2011, and started with a few test crawls. We immediately noticed that Silk Road relies on authentication cookies that can be reused for up to a week without having to re-authenticate through the login prompt of the website. Provided we can manually refresh the authentication cookie at least once per week, this allows us to bypass the CAPTCHA mechanism and automate our crawls.

We conducted a near-comprehensive crawl of the site on November 29, 2011,¹ using HTTrack [34]. Specifically, we crawled all “item,” “user” (i.e., seller) and “category” webpages. The complete crawl completed in about 48 hours and corresponded to approximately 244 MB of data, including 124 MB of images.

Starting on February 3, 2012, and until July 24, 2012, we attempted to perform daily crawls of the website. We noticed that early in 2012, Silk Road had moved to inlining images as `base64` tags in each webpage. This considerably slowed down crawls. Using an incremental mode, that is, ignoring pages that had not changed from one crawl to the next, each of these crawls ran, on average, for about 14 hours. The fastest crawl completed in slightly over 3 hours; the slowest took almost 30 hours, which resulted in the following daily crawl to be canceled. To avoid confusion between the time a crawl started, and the time a specific page was visited, we recorded separate timestamps upon each visit to a given page.

3.2 Challenges

Kanich et al. [15] emphasize the importance of ensuring that the target of a measurement experiment is not aware of the measurement being conducted. Otherwise, the measurement target could modify their behavior, which would taint the measurements. We thus waited for a few days after the November crawl to see if the full crawl had been noticed. Perusing the Silk Road forums [6], we found no mention of the operators noticing us; our account was still valid and no one contacted us to inquire about our browsing activities. We concluded that we either had not been detected, or that the operators did not view our activities as threatening.

We spent some additional effort making our measurements as difficult to detect as possible. Since all Silk Road traffic is anonymized over Tor, there is no risk that our IP address could be blacklisted. However, an identical Tor circuit (on our side) could be repeatedly used if our crawler keeps the same socket open; this in turn could reveal our activities if the Silk Road operators monitor the list of Tor circuits they are running, and realize that a fixed Tor rendez-vous point is constantly being used. We addressed this potential issue by ensuring that all circuits, including active circuits, are periodically discarded and new circuits are built. To further (slightly) obfuscate our activities, instead of always starting at the same time, we started each crawl at a random time between 10pm and 1am UTC.

Despite all of these precautions, we had to discard some of our data. On March 7, 2012 a number of changes were implemented to Silk Road to prevent profiling of the site [6]. Whether this was due to Silk Road operators noticing our crawls or to other activity is unclear. URL structure changed: item and users, instead of being referenced by a linearly increasing numeric identifier, became unique hashes. Fortunately, these hashes initially simply consisted of a substring of the MD5 hash of the numeric identifier, making

¹All dates and times are expressed in Universal Time Coordinates (UTC).

Books

Bookmark this item

Hacking for beginners

Seller:
Hacking Tech (98)

Price:
\$0.12

Ships from: undeclared
Ships to: Worldwide

Description:
Hacking For Beginners is a reference book for beginners to learn ethical hacking for free and from basic level to clear all the fundamental concepts of ethical hacking.the book has been prepared by Hacking Tech (www.hackingtech.co.tv) website for the users benefit.so enjoy the book and site...



add to cart

Recent feedback

rating	feedback	freshness
5 of 5	Fast delivery	3 days
5 of 5	Thanks!	4 days
5 of 5	Leave feedback here	9 days
5 of 5	Leave feedback here	9 days
5 of 5	5 of 5	10 days

Figure 2: Silk Road item page. Each item page contains seller, price, and shipping information, as well as buyer feedback on the item.

it easy to map them to the original identifiers.² More problematically, feedback data, which is crucial to estimating the volume of sales became aggregated and feedback timestamps disappeared. That is, instead of having, for an item G sold by S a list of n feedback messages corresponding to n purchases of G along with the associated timestamps, Silk Road switched to presenting a list of 20 feedback messages, undated, across *all* the items sold by S . In other words, feedback data became completely useless. Thankfully, due to very strong pushback from buyers who argued that per-item feedback was necessary to have confidence in purchases [6], Silk Road operators reverted to timestamped, per-item feedback on March 12, 2012. Nevertheless, we had to discard all feedback data collected between March 7, 2012 and March 12, 2012.

Finally, in several instances, Silk Road went down for maintenance, or authentication was unsuccessful (e.g., because we had not refreshed the authentication cookie in time), leading to a few sporadic days of missing data. The largest gaps are two eight-day gaps between April 10, 2012 and April 17, 2012 due to an accidental suspension of the collection infrastructure; and between July 12, 2012 and July 19,2012, due to an accidental deletion of the authentication cookie.

3.3 Data collected

We only collect data that is both publicly accessible over the Tor network, and linked from other Silk Road pages. That is, we do not collect buyer data, as buyers do not have public “buyer pages.” We also do not collect stealth listings, or data about sellers when they operate in stealth mode.

We primarily focus data collection on “item pages,” that is, pages describing the goods being sold on Silk Road. We show an example in Figure 2. Each item page is bound to a unique item identifier as part of its URL (integer until March 7, 2012, 10-digit hash afterwards), and contains the name of the item (“Hacking for beginners” in Figure 2), a picture, the category in which the item fits (e.g., “Books”), seller information (a name, percentage of positive feedback, and a hyperlink denoting the seller unique identifier), price (e.g., 0.12 BTC), shipping information, item description, and buyer feedback. We gather all of this information for each item we crawl, and record a timestamp (in UNIX epoch time) every time the page is visited.

Feedback data. Each piece of feedback consists of three fields: a rating between 1 and 5, a textual description of the feedback, and the age of the feedback. Feedback age is expressed in minutes, hours, days or months, depending on how old the feedback is. Hence, we can timestamp much more accurately feedback recently given at the time of the crawl, than older feedback. This is one of the reasons for crawling Silk Road daily: the age of feedback less than a day old can be quite precisely pinpointed.

We record feedback in two different manners. For each crawl of Silk Road started at time t and lasting until $t + \tau$ ($\tau > 0$), we record all feedback present on the site in a separate database \mathcal{D}_t , thereby getting a *snapshot* of the feedback amassed until time $t + \tau$. This method may miss some feedback. For instance, if we crawl an item page at time $t + \tau_1$, and a customer leaves feedback at time $t + \tau_2$ with $\tau_1 < \tau_2 < \tau$, that customer’s feedback will not be recorded as part of the time- t snapshot. Furthermore, timestamps of feedback given long before t may be very approximate.

To address this issue, we also record, in a database \mathcal{D} , novel feedback from one crawl to the next, that is, feedback for which text did not previously appear in our records for this specific item. This method guarantees that feedback timestamps are as accurate as possible (since they are recorded as soon as the feedback is observed). Furthermore, we can capture nearly all the feedback present on the site, without worrying about collection gaps. A drawback of this method is that it may overestimate the amount of feedback when there are feedback updates. In particular, new buyers are sometimes asked to finalize early, that is, to send feedback immediately after the online transaction is completed and before receiving goods. They may elect to update the feedback after delivery of the goods purchased, which can be weeks later. When this happens, the original feedback is replaced on the website by the new feedback, and the timestamp is updated. However, \mathcal{D} contains both the original, and the updated feedback(s), even though only one sale occurred.

Maintaining both a family (\mathcal{D}_t) of database of snapshots of the site, and a cumulative database \mathcal{D} allows us to have lower and upper bounds on the amount of feedback posted on the site, which in turn is a useful indicator of sales.

4 Marketplace characteristics

We next provide an overview of the types of goods being sold in Silk Road, before discussing seller characteristics.

²New identifiers subsequently created are salted hashes with a non-trivial salt; but those do not map to items and users that had already been registered on the site when the switch occurred. Thus, we do not need to find the pre-images of these hashes and can instead simply treat them as unique identifiers.

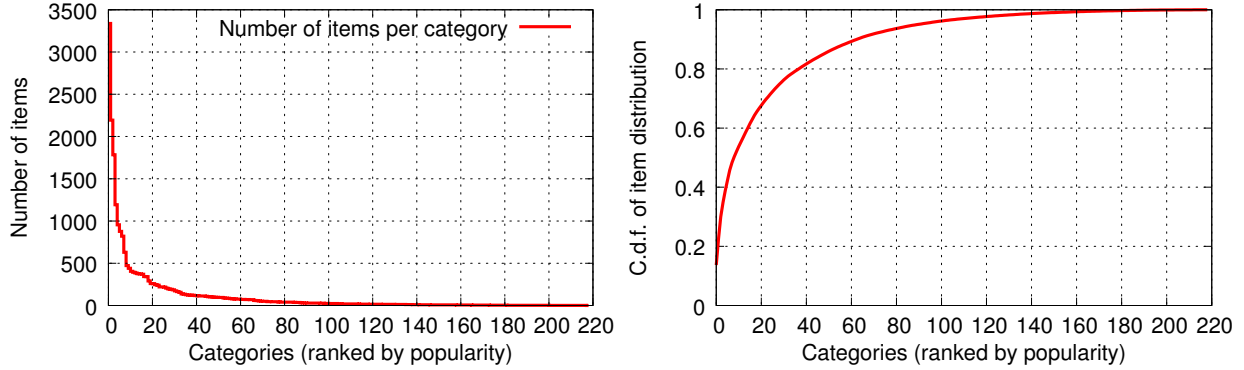


Figure 3: **Distribution of items per category.** The plots show the number of items in each category, ordered by decreasing popularity (left) and the cumulative distribution of all items over all categories (right). The 20 most popular categories represent over 2/3 of all items available.

4.1 What is being sold?

Items offered on Silk Road are grouped by categories. There are approximately 220 distinct categories, ranging from digital goods to various kinds of narcotics or prescription medicine. In Figure 3, we plot, on the left-hand side, for each category, the number of items sold in that category, over the data collected from February 3, 2012 through July 24, 2012. For readability we ordered categories by decreasing popularity. In total, we found 24,385 unique items being sold over that period. While a few categories seem to hold the most items, Silk Road, like other online marketplaces, exhibits a long-tail behavior, where a large number of items appear to be unique. This is confirmed by the right-hand graph, where we plot the cumulative distribution of items as a function of the number of categories considered. The right-hand graph shows that over two thirds of all products sold on Silk Road during our data collection interval belong to one of the top 20 categories, but that, after that, the cumulative fraction of items only slowly converges to 1.

In Table 1, we take a closer look at the top 20 categories per number of item offered. “Weed” (i.e., marijuana) is the most popular item on Silk Road, followed by “Drugs,” which encompass any sort of narcotics or prescription medicine the seller did not want further classified. Prescription drugs, and “Benzos,” colloquial term for benzodiazepines, which include prescription medicines like Valium and other drugs used for insomnia and anxiety treatment, are also highly popular. The four most popular categories are all linked to drugs; nine of the top ten, and sixteen out of the top twenty are drug-related. In other words, Silk Road is mostly a drug store, even though it also caters some other products. Finally, among narcotics, even though such a classification is somewhat arbitrary, Silk Road appears to have more inventory in “soft drugs” (e.g., weed, cannabis, hash, seeds) than “hard drugs” (e.g., opiates); this presumably simply reflects market demand.

Item availability. We estimate item availability by first recording the first time we saw an item being listed, and the last time we saw it listed. Items may have been listed and de-listed several times in the meantime; here we are only looking at the overall lifespan of an item, regardless of its transient availability. To account for the large number of items still available when we stop collection on July 24, 2012, we use a standard Kaplan-Meier estimator [16] to compute the “survival probability,” $S_{\text{item}}(\tau)$, of a given item after τ days, that is, the probability an item is listed for more than τ days. $S_{\text{item}}(\tau)$ is plotted in Figure 4 along with 95% confidence intervals.

We discover that a majority of items disappear within three weeks of their first being listed ($S_{\text{item}}(21) = 0.473$, 95% c.i. (0.466, 0.479)); and more than 25% of items disappear within three days ($S_{\text{item}}(3) =$

Category	#. items	Pct.
Weed	3338	13.7%
Drugs	2194	9.0%
Prescription	1784	7.3%
Benzos	1193	4.9%
Books	955	3.9%
Cannabis	877	3.6%
Hash	820	3.4%
Cocaine	630	2.6%
Pills	473	1.9%
Blotter (LSD)	440	1.8%
Money	405	1.7%
MDMA (ecstasy)	393	1.6%
Erotica	385	1.6%
Steroids, PEDs	376	1.5%
Seeds	374	1.5%
Heroin	370	1.5%
DMT	343	1.4%
Opioids	342	1.4%
Stimulants	291	1.2%
Digital goods	260	1.1%

Table 1: **Top 20 categories in terms of items available.** Products sold on Silk Road are mostly listed as narcotics or controlled substances.

0.745, 95% c.i. (0.739, 0.750)). On the other hand, there are also a few very long-lived items (on the right-hand side of the graph) that have been present for the entire collection interval. There may be two different explanations for the relatively short lifespan of each item: vendors may run out of stock quickly and de-list their items, possibly re-listing them later under a slightly different name resulting in a different item page, or they may elect to make them stealth listings as soon as they have established a customer base.

Custom listings. Finally, public custom listings are relatively rare. Out of the 24,385 items we observed, only 745 were explicitly marked as “custom listings.” This is undoubtedly a lower bound, as custom listings should be stealth listings, except when sellers are running tests.

4.2 Who is selling?

Due to the anonymous nature of Silk Road, it is impossible to discern whether certain sellers use multiple seller pages; economically, this is not an attractive proposition, as there is a fee associated with opening a seller account. Likewise, several sellers in the physical world may offer their goods through a unique seller page on Silk Road, although this would certainly be a clear indication of a business partnership. In this discussion we will equate “sellers” with distinct seller pages.

Evolution over time. In Figure 5, we plot the evolution of the number of sellers on Silk Road over time, between February 3, 2012 until our last daily crawl (July 24, 2012).

On July 24, 2012, we found 564 distinct sellers with at least one item listed for sale on Silk Road. This

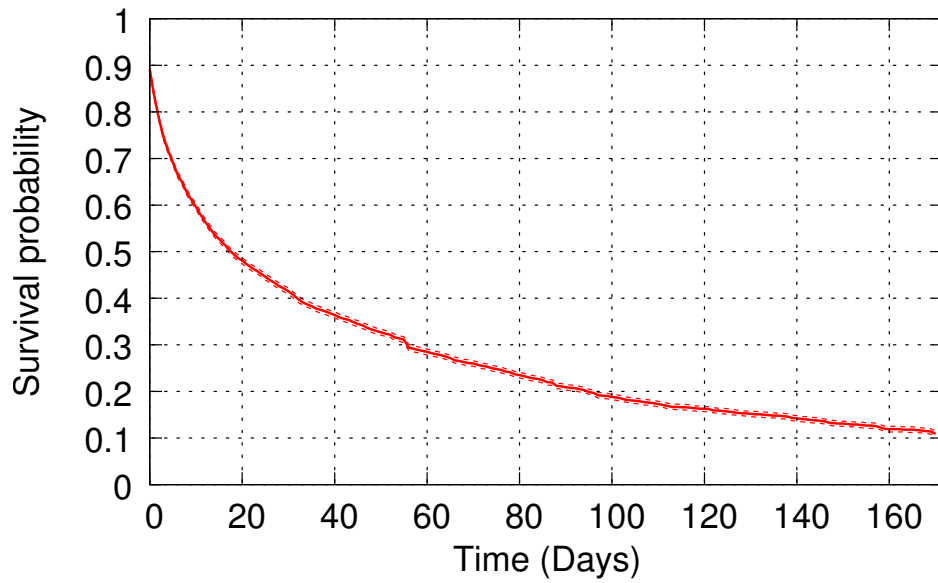


Figure 4: **Probability that a given item will be available on the site as a function of time (in days).** The dotted lines correspond to the 95% confidence interval. Most items are only available for limited periods of time, with a majority of items being available for less than three weeks.

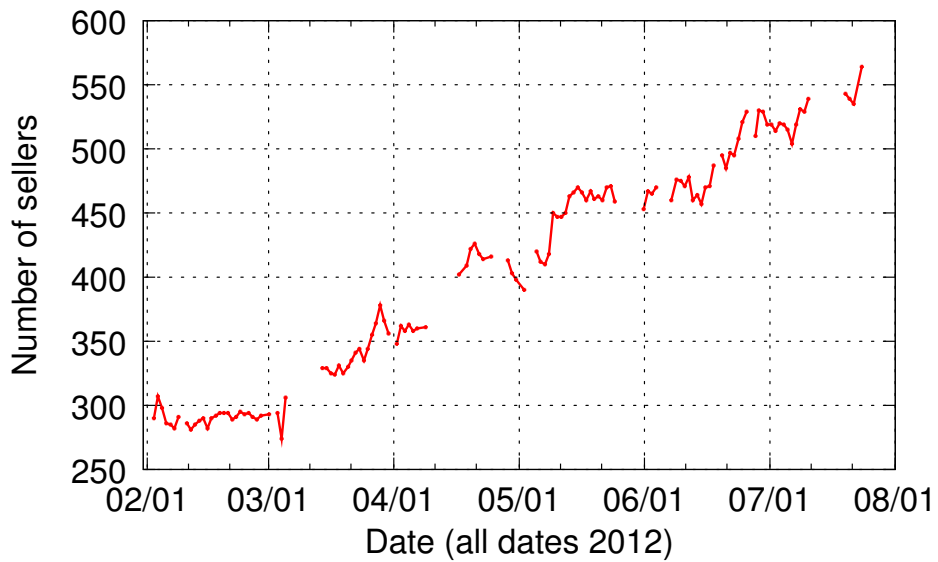


Figure 5: **Evolution of the number of sellers.** As a point of comparison, there were only 220 sellers on Nov. 29, 2011.

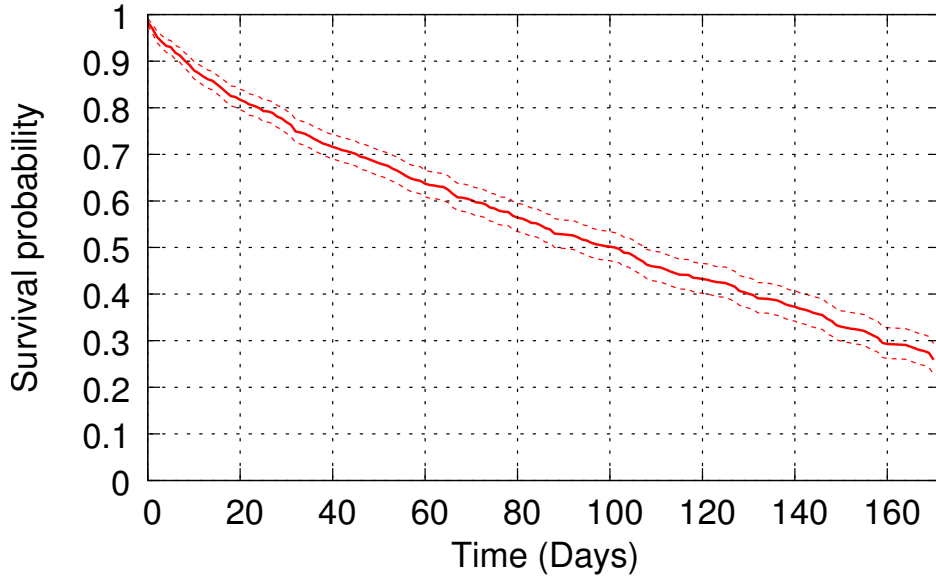


Figure 6: **Probability a seller will continue publicly selling items on the site as a function of time (in days).** The dotted lines denote the 95% confidence interval. About half of the sellers leave the site within 100 days of initial appearance; around a fifth of the sellers stay for less than three weeks.

is a marked increase compared to the 220 sellers we had observed during our initial crawl on November 29, 2011 (not shown on the figure). The gaps in the figure correspond to data collection gaps. An interesting spike occurs around April 20, 2012. April 20 featured a large promotional sale on Silk Road to mark “Pot Day.” It appears that a number of sellers entered the marketplace in the week or two prior to this operation; and a non-negligible number left immediately afterwards. Also, the Silk Road forums indicate that one of the top sellers went on hiatus on March 12, 2012. It is unclear whether this played a role in the marked increase of the number of sellers since that time – i.e., whether newcomers attempted to fill the void. The number of active sellers has been continuously increasing, at least since early March. A linear regression fit in Figure 5 gives $y = 1.674x + 265.605$ ($R^2 = 0.969$), where x is in days, y the number of active sellers, and the time origin is set to the beginning of our measurement interval. In other words, the increase in the number of sellers appears linear, with about 50 new active sellers each month.

However, this simple regression analysis does not accurately reflect the fact that many sellers actually leave the marketplace. Over the measurement interval February 3, 2012– July 24, 2012, we indeed found 1,239 distinct sellers were at one point or another publicly listed. We use again a Kaplan-Meier estimator to plot the probability $S_{\text{seller}}(\tau)$ that a given seller will remain active on the site for at least τ days. Similar to our discussion about item availability, we consider a seller is active by subtracting the first time they were seen from the last time they were seen, regardless of their coming and going in between. Figure 6 plots the function S_{seller} , along with 95% confidence intervals. Compared to item turnover, seller turnover is rather modest. A majority of sellers stay on the site for approximately 100 days ($S_{\text{seller}}(101) = 0.500$, 95% c.i. (0.469, 0.532)); about a fifth of all sellers are present for less than three weeks ($S_{\text{seller}}(21) = 0.813$, 95% c.i. (0.791, 0.836)). A “core” of 112 sellers (9% of all sellers) were present for the entire measurement interval. All of this data point to relative seller stability.

Geographic location. We next inspect the advertised origins and acceptable shipping destinations for each

Origin		Acceptable destinations	
Country	Pct.	Country/Region	Pct.
U.S.A.	43.83%	Worldwide	49.67%
Undeclared	16.29%	U.S.A.	35.15%
U.K.	10.15%	European Union	6.19%
Netherlands	6.52%	Canada	6.05%
Canada	5.89%	U.K.	3.66%
Germany	4.51%	Australia	2.87%
Australia	3.19%	World. except. U.S.A.	1.39%
India	1.23%	Germany	1.03%
Italy	1.03%	Norway	0.70%
China	0.98%	Switzerland	0.62%
Spain	0.94%	New Zealand	0.56%
France	0.82%	Undeclared	0.26%

Table 2: **Top 12 most frequent shipping origins (left), and acceptable shipping destinations (right).** Certain sellers ship to multiple destinations, hence totals may exceed 100%.

of the 24,385 items. Table 2 shows the top 12 locations for both origin and destinations. Some items ship to multiple destinations (e.g., Norway, Switzerland and European Union) so that the total of the rightmost column adds up to more than 100%. Most items ship from the United States, with the United Kingdom a distant second. The Netherlands are also strongly represented, which is not necessarily surprising given the relatively permissive narcotics laws in the country. We note a clear bias toward English-speaking countries which represent almost two-thirds of all listed origins. This is not surprising since all communications on Silk Road are in English.

More surprisingly, we note that a majority of items ship worldwide, in spite of the nature of the items, as discussed above. One would think that sellers may be reluctant to ship narcotics across borders. It turns out not to be the case, for a couple of reasons. First, sellers with an established reputation may often demand that buyers pay upon purchase, and before delivery of the item. If the item is not delivered, the buyer may have very little recourse, particularly if they have not established a strong reputation in the marketplace. While, as discussed in Section 2, Silk Road offers an escrow service, disputes arising after “early finalization” are considerably harder to mediate by the marketplace operators. Second, the quantities being sold are generally rather small (e.g., a few grams of marijuana), and tracing the senders may be a very difficult task as they can use couriers to mail the items rather than going to the post-office themselves. Third, most sellers use techniques to make package inspection unlikely – e.g., vacuum sealing, “professional-looking” envelopes with typed destination addresses [5]. In other words, sellers can expand their customer base at a relatively low risk for them. Economic incentives justify worldwide shipping, especially since sellers can factor in their selling price the risk of package seizure, and accordingly offer loyal customers at least partial refund guarantees.

Seller ranking. As discussed above, we observed 1,239 unique seller accounts on Silk Road between February 3, 2012 and July 24, 2012, with between 281 and 564 active at a time. Figure 7 shows the proportion of items in the marketplace as a function of the number of sellers. No seller is selling a large inventory. Instead, each seller accounts, at most, for 1.5% of the total number of items found on Silk Road.

While it does not appear like a given seller (or group of sellers) sells a significant proportion of items

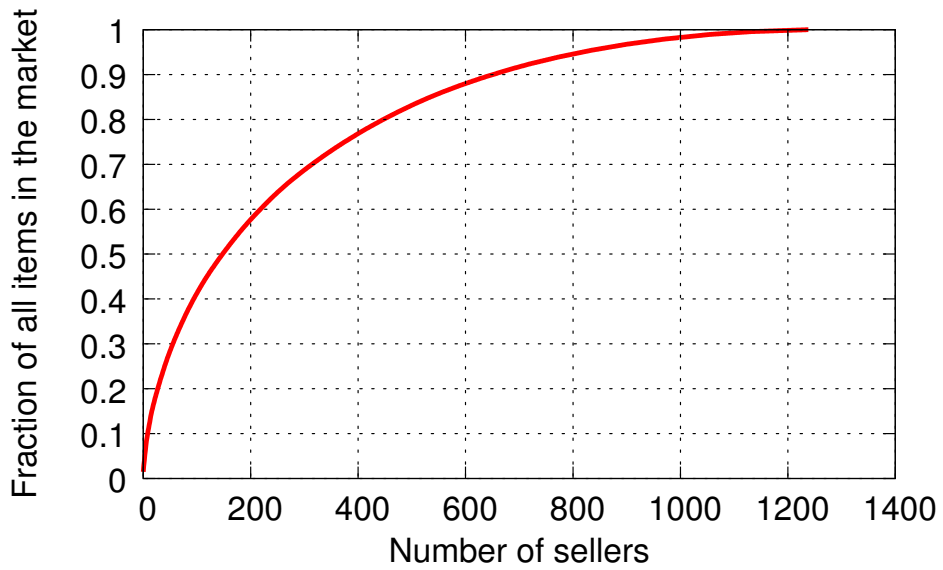


Figure 7: **Proportion of items in the marketplace as a function of the number of sellers.** We observe fairly high diversity, with each seller selling at most 1.5% of the total number of items in Silk Road.

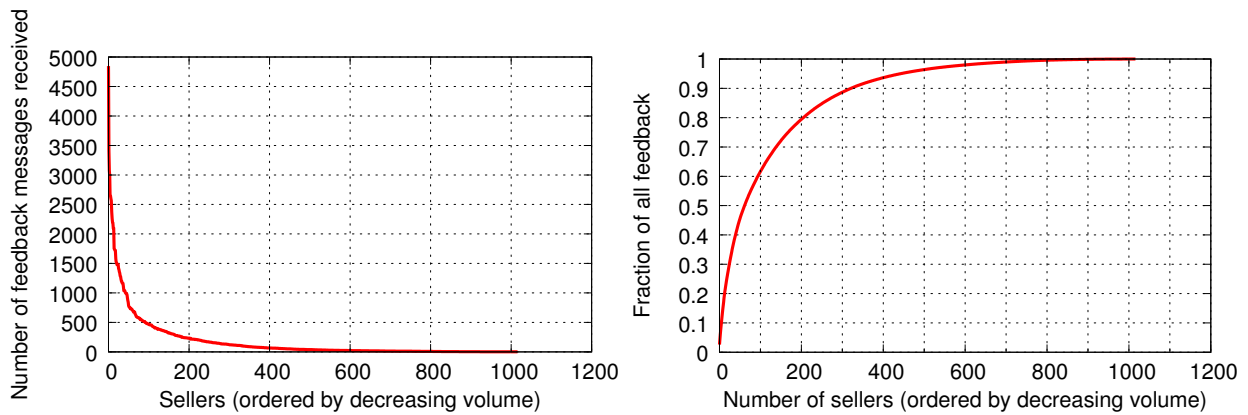


Figure 8: **Number of feedback received per seller (left) and cumulative fraction of all feedback observed over the collection interval as a function of the number of sellers (right).** The top 100 sellers are responsible for approximately 60% of all feedback gathered.

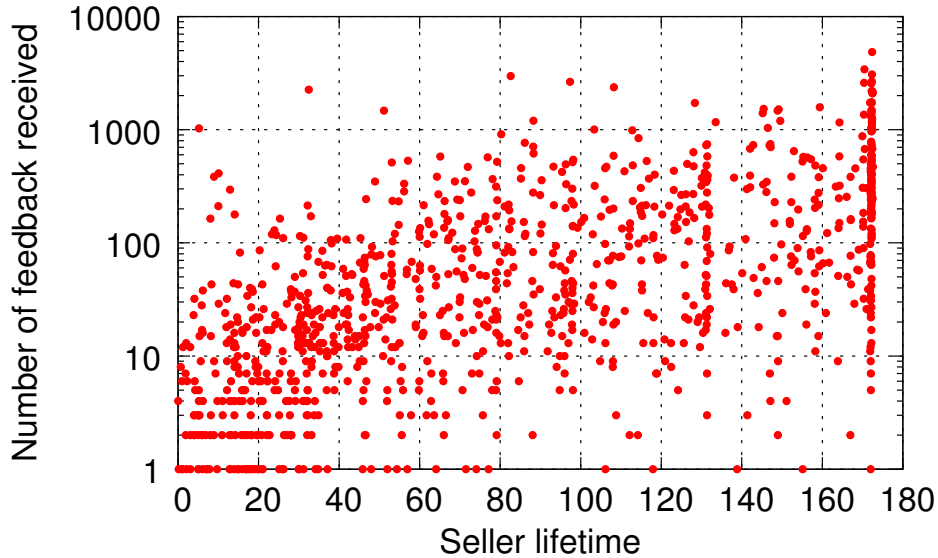


Figure 9: **Amount of feedback received by sellers vs. the amount of time sellers have been listed on the site.** There is a modest positive correlation between the two quantities with Pearson’s correlation coefficient $r \approx 0.39$.

overall, it could be the case that a few selected items sell in large quantities. As discussed in Section 3 we use the amount of feedback collected as a proxy for the number of sales made. In Figure 8, we plot, on the left-hand side, the number of feedback received per seller (where sellers are ranked in decreasing amount of feedback received), and on the right-hand side the cumulative fraction of total feedback as a function of the number of sellers considered. These plots show that a few sellers indeed receive a large amount of feedback, in absolute terms. For instance, the seller with the largest volume of sales received 4,847 feedback messages over the six months we monitored. However, the market is quite spread out between sellers, as can be seen on the right-hand graph. Roughly 100 sellers correspond to 60% of all feedback gathered.

Considering we previously observed that 112 sellers have been present throughout our measurements, we hypothesized that this core was perhaps the most active. In Figure 9, we use a scatter-plot to graph the amount of feedback received (in logarithmic scale) by each seller against the number of days the seller has been present on the site. While the graph suggests a positive correlation between the two quantities, that correlation appears quite weak. This is confirmed by computing Pearson’s correlation coefficient between the two datasets: we get $r \approx 0.39$. In other words, there are a number of sellers that do not appear to stay very long on the site, but amass a large number of transactions, while on the other hand, some of the “old timers” appear to relatively rarely engage in transactions. These results may be due to the sellers electing to go in stealth mode after having built up a large enough user base, or to list most of their products as stealth listings.

4.3 Customer satisfaction

We next discuss customer satisfaction. On a site like Silk Road, where seller anonymity is guaranteed, and no legal recourse exists against scammers, one would expect a certain amount of deception. Most transactions seem, however, to generate excellent feedback from buyers. Table 3 provides a breakdown of the feedback

Rating	Number	Pct.
5/5	178341	96.5%
4/5	2442	1.3%
3/5	1447	0.8%
2/5	520	0.3%
1/5	2053	1.1%

Table 3: **Distribution of feedback ratings.** A vast majority of transactions seems to proceed to the satisfaction of the buyers.

ratings from 184,804 feedback instances we collected. 97.8% of feedback posted was positive (4 or 5 on a scale of 1 to 5). In contrast, only 1.4% of feedback was negative (1 or 2 on the same scale).

Thus, it appears at first glance that Silk Road sellers are highly reliable; or, at the very least, that the escrow system used is effective at policing the marketplace. We caution, however, against too rapid an interpretation of this result. First, a study by Resnick and Zeckhauser [33] shows that Internet users in general disproportionately use positive feedback when rating online experiences. In fact, over 99% of the feedback in the eBay corpus used in their study [33] was positive. Second, not all transactions have feedback reported. Indeed, a number of transactions are made “out of escrow,” i.e., directly between a seller and a buyer. For those, there is no feedback mechanism, nor any oversight possible from the Silk Road operators. We suspect most of the scams³ occur “out of escrow,” and of course, no feedback for these transactions is reported since they technically do not exist (from the market operator standpoint).

Finalizing early. We observe that 20,884 instances of feedback contain variations of “F.E.,” or “finalizing early,” accounting for spelling variations (“finalize” vs. “finalise”) and word order (“early finalization” vs “finalize early”). This shows that finalizing early is a rather common practice on Silk Road. There does not appear to be significantly more problems reported with feedback including such strings (only 342 of them map to a rating of 1 or 2). This seems to show that established sellers that are offered the option of requesting early finalization from their customers do not abuse that privilege.

5 Economic aspects

We next present a brief discussion of economic indicators on Silk Road. We start by looking at the evolution of the prices of a basket of items. We then turn to estimating transaction volumes that occur on Silk Road.

5.1 Inflation

All transactions on Silk Road are using Bitcoins (BTC). Bitcoin has been a notoriously volatile currency, going from 1 BTC being worth around 30 cents in January 2011 to 1 BTC reaching USD 31.90 in an intra-day high on June 8, 2011, and declining rapidly back to approximately 1 BTC≈USD 2.5 in late October 2011.

In Figure 10, we plot the evolution of the exchange rate of the Bitcoin against the three major currencies that sellers use in their countries, over the duration of our measurements. As can be seen in the figure, the Bitcoin exchange rate has remained relatively stable between the end of February and early May, oscillating

³The Silk Road operators in fact warn that buyers relying on out of escrow transactions “have been scammed [5].”

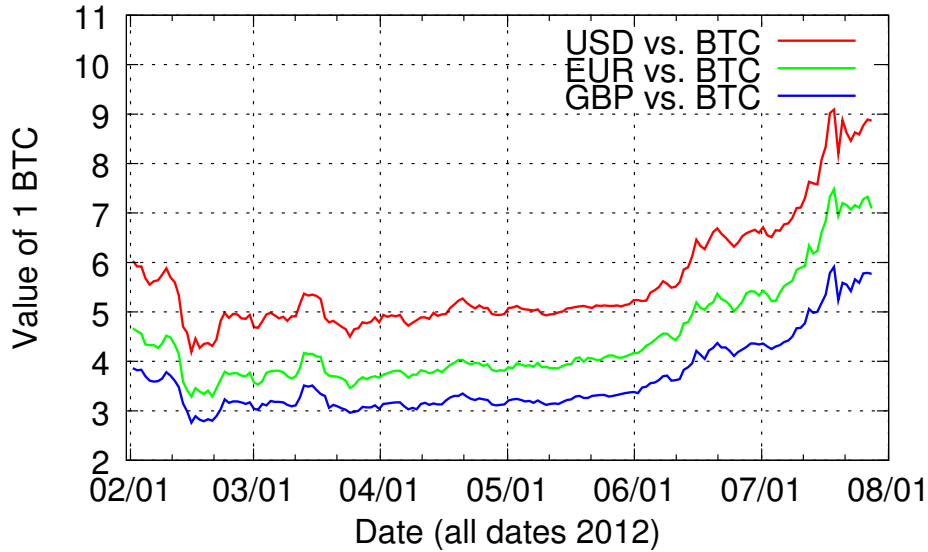


Figure 10: **Evolution of the value of a Bitcoin in the three major currencies used in the countries of the sellers operating on Silk Road, over our collection interval.** Each point corresponds to a weighted average over a given day. Data is from Mt.Gox [4].

around $1 \text{ BTC} \approx \text{USD } 5$, and corresponding values in euros and pounds. Since then, the Bitcoin has notably appreciated, reaching close to USD 9 since mid-July 2012, with relatively large fluctuations.

We look at the evolution of the normalized price of the five best-selling items on Silk Road, corresponding to between 1,025 and 1,590 pieces of feedback collected. These five items are all narcotics, but correspond to four different categories, and have prices ranging from 1 BTC to approximately 50 BTC. For any time t , the normalized price compared to origin t_0 is defined as $P(t)/P(t_0)$.

We plot the normalized prices of these five items in Figure 11. On the left, we observe that they quite closely mirror the fluctuations of the Bitcoin exchange rate: as the Bitcoin appreciates, the prices drop; conversely, a drop in the Bitcoin value results in a price increase.

Silk Road provides automatic pegging of prices to the US dollar if sellers so desire. On the right, we plot the normalized price in US dollars. The plot confirms that there is generally no evidence of rapid inflation (or deflation) on Silk Road as all prices remain roughly within $\pm 10\%$ of their original price. We notice a drop in price prior to April 20 for all items. Indeed, this price decrease appears to be part of a promotional campaign for “Pot Day.” A second observation is that item 2 stops being sold immediately after April 20. The last time it is observed on the site is April 25, before being de-listed. From discussions in Silk Road forums [6], it appears that the seller of that item abruptly left the marketplace, potentially leaving a large number of paid, finalized early, orders unfulfilled. In other words, there is suspicion of a “whitewashing attack [12],” whereby a seller creates an excellent reputation, before using that reputation to defraud users and leave the system. In hindsight, the 20% drop in price occurring just prior to April 20 was considerably steeper than all the other promotional discounts. This could have been an indicator that the seller was not intending on fulfilling their orders and was instead artificially lowering prices in hopes of attracting large numbers of customers to defraud.

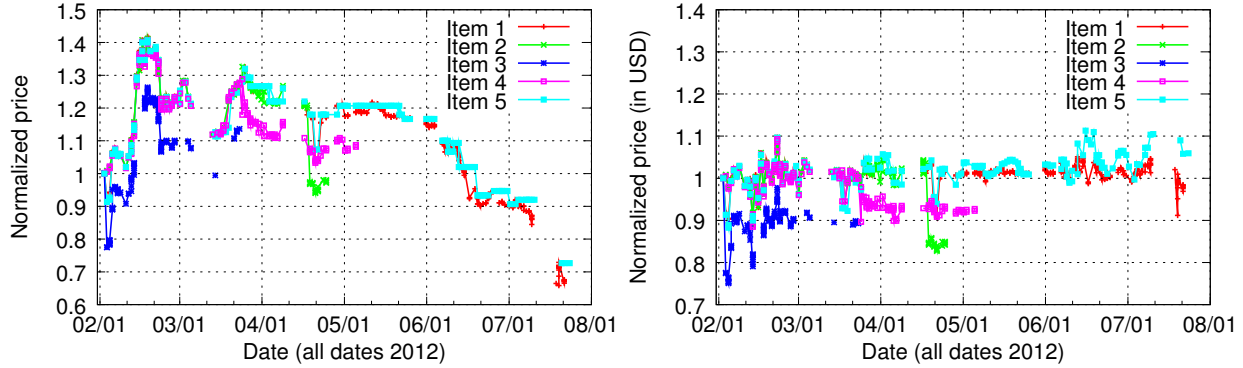


Figure 11: **Evolution of the normalized price of the five most sold items on Silk Road in BTC (left) and US dollar (right).** The evolution closely mirrors the evolution of the Bitcoin exchange rates, suggesting little inflation for these items over the time interval considered.

5.2 Transaction volumes

We next provide an estimate of the total amount of daily sales realized in Silk Road. Obtaining this estimate is problematic, because, as explained in Section 3, feedback data is relatively noisy. Feedback can be updated, which results in a timestamp update, and in discarding the old feedback; and feedback is not always issued at the time the item is purchased, but instead, when it is received.

We decided to use averages over a sliding window of size 29 days to try to assess the daily sales volumes. The averaging mechanism smoothes out potential issues due to delayed feedback. The reason why we limit ourselves to 29 days is that timestamps on feedback older than 30 days (i.e., one month) are extremely approximate, with potential errors of 29 days in the worst case, while timestamps on feedback less than a month old are accurate within 24 hours.

We present our results in Figure 12. Each point corresponds to the average daily sales volumes over the past 29 days.⁴ The sales volumes are computed by, for each item, counting the number of feedback gathered about the item over the past 29 days, and multiplying this number by the average price of the item over the last 29 days. The estimated daily sales volumes are then simply taken as the resulting sales volume divided by 29.

We observe that the total volume of sales has been increasing quite significantly, going from approximately 6,000 BTC/day to approximately 9,500 BTC/day, before seemingly retreating down to 7,000 BTC/day.

The latter decrease is, however, an artifact of the Bitcoin sharply appreciating against all major currencies, rather than an indication of a drop in sales. Over our entire collection interval, the daily volume of sales approaches 7,665 BTC/day.

Converting to US dollars, using the Bitcoin exchange rate plotted in Figure 10, we obtain a total sales volume of over USD 1.22 million per month when averaged over our measurement interval. This would correspond to an annual revenue of close to 15 million USD for the entire marketplace.

⁴We use a 29-day as opposed to a 30-day window, as 30-day old feedback may be marked as “30 days old” or as “one month old” depending on the exact time when it is collected, resulting in significant errors in feedback counts.

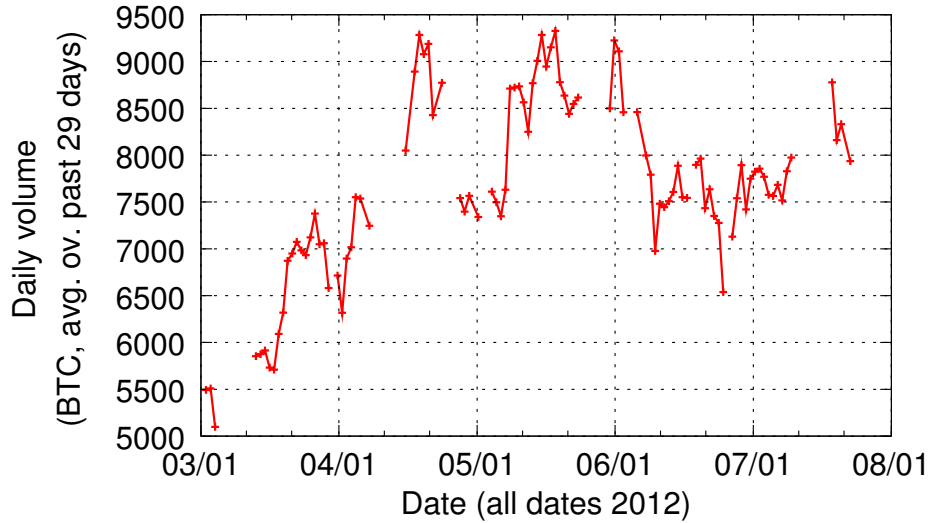


Figure 12: **Estimate of the total amount of daily sales (in BTC) occurring on Silk Road.** Each point corresponds to an average over the prior 29 days.

Item price	Op. commission
first \$50	10%
\$50.01 - \$150	8.5%
\$150.01 - \$300	6%
\$300.01 - \$500	3%
\$500.01 - \$1000	2%
over \$1000	1.5%

Table 4: **Silk Road operator commission schedule (from January 2012).** Prior to the establishment of this schedule, the commission was at a flat 6.23% rate.

5.3 Operator commissions

Silk Road operators collect a commission on all sales realized on the website. The commission schedule was originally set at 6.23% of the sales price. In January 2012, a tiered commission schedule was established, using a model reminiscent of eBay’s fee structure. The schedule, described in Table 4, is based on amounts in US Dollars, not in Bitcoins. The first USD 50 of a sale are charged a 10% commission fee, then, from USD 50.01 to USD 150, the fee becomes 8.5%, and further decreases as the value of the item sold increases.

In Figure 13, we plot an estimate of the daily commissions collected by Silk Road operators as a function of time. We simply reuse the previous estimates, and apply both the fixed 6.23% rate, and the schedule of Table 4 to each item. We find that the new schedule turns out to yield on average a commission corresponding to approximately 7.4% of the item price.

Using the new schedule we find that Silk Road operators have seen their commissions increase from over USD 2,200/day in March 2012 to roughly USD 4,000/day in late July 2012. In other words, even though the volume in Bitcoins may have decreased due to the Bitcoin rising against the US dollar, the

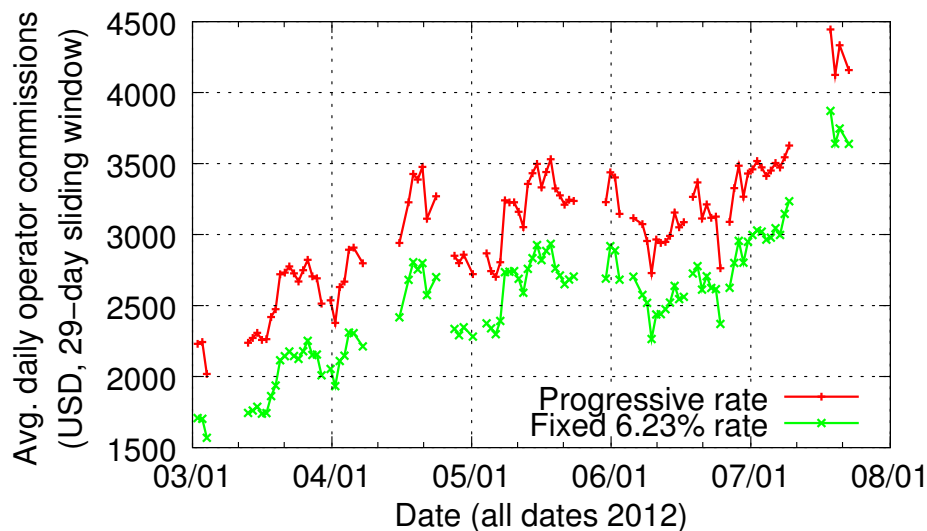


Figure 13: **Estimate of the daily commission collected by Silk Road operators (in USD).** Each point corresponds to an average over the prior 29 days. The progressive rate schedule is given in Table 4.

transaction volume in US dollars, and the corresponding commissions, have significantly increased over our measurement interval. Over the entire measurement interval, we compute that Silk Road operators collect, an average, roughly 92,000 USD per month in commissions. Stated differently, projected over a year, Silk Road operators’ revenue would probably be around 1.1 million US dollars.

5.4 Silk Road in the Bitcoin economy

Finally, we were interested in computing the share of Silk Road trade in the overall Bitcoin economy. This is a notoriously difficult quantity to estimate, as quantifying the total volume of Bitcoin transactions is, itself, challenging. In particular, the use of tumblers and mixers implies that each Silk Road transaction corresponds to a large number of Bitcoin transactions.

As a potential indicator, we suggest to compare the estimated total volume of Silk Road transactions with the estimated total volume of transactions at all Bitcoin exchanges (including Mt.Gox [4], but not limited to it). The latter corresponds to the amount of money entering and leaving the Bitcoin network, and statistics for it are readily available [2].

Over our entire measurement interval, using the 29-day moving averages earlier described, approximately 1,335,580 BTC were exchanged on Silk Road. Likewise, computing 29-day moving averages for all transactions recorded in Bitcoin exchanges, approximately 29,553,384 BTC were traded in Bitcoin exchanges over the same period. This number is larger than the entire supply of currency, since each coin is exchanged several times. Comparing the two numbers shows that Silk Road transactions correspond to about 4.5% of all transactions occurring in exchanges.

This estimate is however not particularly robust. Indeed, for each item purchased on Silk Road, the buyer has to procure, directly or indirectly, Bitcoins from an exchange,⁵ and eventually the seller may want

⁵Buyers could have also mined Bitcoins, but we hypothesize this is comparatively rare.

to redeem funds in their currency, which corresponds to two exchange transactions. Alternatively, sellers may directly re-invest Bitcoins, e.g., to purchase other items on Silk Road, in which case at most one exchange transaction occurs.

The only conclusion we can draw from this comparison is that Silk Road-related trades could plausibly correspond to 4.5% to 9% of all exchange trades. While far from being a negligible share of the Bitcoin economy, speculative trades (i.e., using Bitcoin as a commodity rather than a currency) still constitute the bulk of all exchange trades. A much more thorough analysis would be required to properly assess the various components of the Bitcoin economy, but is outside the scope of this paper.

6 Discussion

While we believe this research may bring up a number of discussions, and hopefully even start a public policy debate on the effectiveness of current intervention or prevention strategies for controlled substance abuse, we here choose to narrow our focus to three areas. First, we discuss the accuracy of our estimates. Second we provide an overview of the ethical considerations and associated conclusions we came to during the design of this study. Third, we briefly evaluate potential intervention policies.

6.1 Measurement accuracy of sales volumes

Sales volumes described in Section 5 are obtained through indirect indicators (buyer feedback), and as such only represent an estimate, and not actual, verified volumes.

First, dishonest sellers could plausibly create a number of buyer accounts, leave fabricated feedback to enhance their reputation, and cause our estimates to be inflated. Such behavior, however, is usually relatively easy to spot, and would be reported in Silk Road user forums dedicated to reviewing sellers. From casual observation, the risk of ruining one's reputation for little gain appears to be enough of a deterrent that fabricated feedback only represents a small fraction of all feedback and does not impact our measurements.

On the other hand, buyers typically leave only one piece of feedback *per order*. When an order contains a large quantity of a given item, we are underestimating the total volume of sales taking place. Because this cannot be detected, feedback reports only provide a conservative estimate of overall sales volume. We also cannot account for sales coming from stealth listings. Thus, it is likely that the numbers presented in Section 5 are a lower bound.

6.2 Ethical considerations

Conducting this research yielded some ethical quandaries. Since we are analyzing data from activity that is, in most jurisdictions, deemed as criminal, could this work directly lead to arrest or prosecution of individuals? If such were the case, should it be published?

We answered negatively to the first question. Indeed, the data we collected is essentially public. We did have to create an account on Silk Road to access it; but registration is open to anybody who connects to the site. We did not compromise the site in any way. Perhaps, bypassing the authentication mechanism and associated CAPTCHA by reusing an authentication cookie could be construed as a "hack." However, we argue this is nothing more than using a convenient feature that the site operators have willingly offered their visitors. Indeed, nothing would prevent the site operators from setting authentication cookies with very short expiry dates.

Considering that the data we obtained is available to anybody, we do not think this work adds any additional risk for the Silk Road operators, their customers or their sellers. In fact, as routinely expressed in

user forums [6], Silk Road operators seem to espouse crypto-anarchist ideals (similar to, for instance, those described by May [19]) and to that end, willingly make their website – and, as a result, its data – publicly available.

Research reproducibility. At the same time, we need to allow others to reproduce our results. To that effect, we make a subset of the databases we constructed from our data collection available at <https://arima.cylab.cmu.edu/sr/>. We decided not to make available any textual information (item name, description, or feedback text) because we could not manually inspect each entry to ensure that no potentially private information (e.g., URLs, email addresses) would be released. Among the results presented in this paper, only the numbers on early finalization or stealth listings require textual fields. All other results should be reproducible with the publicly-released dataset.

Using the Tor network for measurement research. Another ethical consideration is linked to the design of the study itself. To acquire the measurements we needed to obtain for analysis, we had to repeatedly crawl the Silk Road website. This in turn resulted in extensively using the Tor network for the purpose of this work. Because Tor is a relatively resource-constrained network, our measurements could have impacted other users that would need it. We believe that the scientific value of this study, and its potential public policy impact justified the use of the network we have made. Monitoring our usage, we realized that we were downloading between 500 and 800 MB of data per day over the Tor network for this project. While certainly not negligible, this remains in the order of a single typical movie download.

However, we do think our usage of the network (for this project, and related research efforts) should be compensated. Partly for this reason, we have recently deployed a fast Tor relay in our institution’s network.

6.3 Potential intervention strategies

Given the nature of the goods sold on Silk Road, it is quite clear that various law enforcement agencies have a strong interest in disrupting Silk Road operations. They appear, so far, to have been unsuccessful since the site is still up and has grown in size since Sen. Schumer called on the U.S. Attorney General and the head of the U.S. Drug Enforcement Agency to put an end to it.

We discuss four possible intervention strategies that could be considered: disrupting the network, disrupting the financial infrastructure, disrupting the delivery model, and laissez-faire.

Attacking the network. The first possible intervention policy is to disrupt the Tor network. Indeed, without Tor, Silk Road cannot operate. This strategy is very likely to be difficult to put in place. First, Tor has many uses beneficial to society – Silk Road and other anonymous online marketplaces are far from representing the majority of Tor traffic, even though this work argues that their importance is growing. Tor is routinely used by oppressed individuals to communicate without fear of reprisal. Thus, disrupting the entire Tor network for the purpose of taking down Silk Road would come at a high collateral cost.

Furthermore, Tor has shown to be resilient to a large number of attacks, due to its open design and to the large amount of academic research it fosters. In particular, Tor hidden services, like Silk Road, have been the subject of considerable scrutiny [29, 31]. Øverlier and Syverson showed that timing and intersection attacks could be used to reveal the location of hidden services. Most of these concerns have been addressed in recent versions of Tor, e.g., through the use of persistent “entry guards.” Murdoch described how covert channels (specifically, clock skew) could leak information allowing to roughly estimate the location of a hidden service.

While a determined adversary, given enough measurements, may be able to roughly estimate the location of a hidden service, pinpointing its exact location, and then being able to prove that the machine is in fact

hosting the hidden content is considerably more challenging if, as we expect, operators of the hidden service suspect they are being monitored [35]. For instance, the hidden service could merely act as a proxy to another machine hosted somewhere else.

Attacking the financial infrastructure. Another possible disruption strategy is to attack the financial infrastructure supporting Silk Road. Bitcoin has shown, in the past, to be a very volatile currency. The June 2011 theft of a large number of Bitcoins from the Mt.Gox exchange [4] actually caused an abrupt collapse of the currency. Certain users have been complaining in forums of the uncertainty on the prices they end up paying due to the instability of Bitcoin and the various commissions they have to pay to purchase Bitcoins, and then to purchase items on Silk Road [6].

Thus, an adversary could also attempt to manipulate the currency to create rapid fluctuations and impede transactions. Besides the obvious collateral costs associated with such strategies, Silk Road does provide hedging mechanisms against short-term fluctuations of Bitcoin. These mechanisms have proven to be enough to allow Silk Road to prosper despite Bitcoin's high volatility – but it is unclear how they would fare in the face of a determined attacker with large monetary resources.

Recent research [32] has also shown that Bitcoin transactions are partially vulnerable to traffic analysis. Indeed, the history of all transactions is publicly available and network analysis can help map sets of public keys to individual users and transactions. Since currency exchanges like Mt.Gox where users redeem Bitcoins for cash bind public keys to actual identities, Bitcoin anonymity guarantees are weaker than most Silk Road users seem to assume, even though additional intermediaries (tumblers) are in place. In particular, large Silk Road sellers withdrawing massive amounts of Bitcoins at once may be relatively easily identified, unless they take additional precautions to hide their tracks.

Attacking the delivery model. Another possible angle of action is to attack the delivery model. That is, to reinforce controls at the post office and/or at customs to prevent illicit items from being delivered to their destination. One interesting finding from is that a large number of sellers seem not to worry about seizures: Most items are marked as shipping internationally, which means that the risk of package loss or destruction is viewed as minimal by the sellers. This is certainly an area that warrants further investigation. In the United States, coordination between agencies is paramount: Customs (which can inspect mail) need to work in concert with Drug Enforcement Agency (DEA) and/or Food and Drugs Administration (FDA), depending on the type of item concerned. Very often, seized packages are simply destroyed, or returned to the sender.

Laissez-faire. Finally, a last possible intervention strategy is actually not to intervene. Politically, this is a questionable proposition, as it may sound as an admission of weakness. There are however studies that show that drug abuse prevention is considerably more cost-efficient than enforcing drug prohibition [9]. From a public support standpoint, recent laws in Colorado and Washington state allow the use of marijuana in these states. Despite the current incompatibility between these laws and federal regulations, they are certainly a sign of an evolution in the public attitude toward drug policies. From an economics standpoint, the relatively rapidly expanding business of online anonymous markets such as Silk Road, and the logistic difficulties in shutting down such markets may further tilt the balance toward prevention and cure.

As a result, laissez-faire, however untenable it might currently appear from a policy standpoint, might become more attractive in light of budget constraints. Although there is no public statement about it, this could be the strategy currently adopted by law enforcement, seeing that the marketplace has not met any significant disruption to its operations, other than transient technical issues, in the nine months we studied it, while at the same time sales volumes have nearly doubled.

7 Related work

From a technical standpoint, this work is closely related to rapidly growing literature on measuring cybercrime (e.g., [13, 17, 18, 20, 22, 24, 25, 26, 27, 36]). The techniques used in this paper (periodic crawls, use of anonymous networks) to collect measurements indeed are relatively common to most work in this field. The main difference between this work and the related cybercrime literature is the object of the measurements. Instead of trying to characterize a security attack or the behavior of an attacker, we are instead trying to describe as precisely as possible an online marketplace. In that respect, our work shares some similarities with works that have tried to model transactions on eBay [14, 33] or Amazon [28]. However, we do not focus much on customer reviews to assess seller reputation but instead primarily use feedback as proxy for sales volume.

At first glance quite similar to our work, McCoy et al. provided a characterization of traffic using the Tor network by monitoring a Tor exit node [20]. Different from this research, we do not actually monitor Tor traffic and instead analyze data posted to an online marketplace. Motoyama et al. [27] performed related measurements to evidence the existence of online “mule” recruitment schemes in crowdsourcing marketplaces.

A more recent paper [22] uses leaked transaction databases to precisely estimate the revenues and profits of three major illicit online pharmacy affiliate networks. Contrary to Silk Road, these networks heavily resort to spam [18] and search-engine manipulation [17] for advertising. McCoy et al. find that, between 2007 and 2011 the gross revenues of each of these illicit affiliate networks range from USD 12.8 million/year to USD 67.7 million/year [22]. Without any advertising other than word-of-mouth, Silk Road, with its USD 15 million/year transaction volume, appears to be comparable in size with these illicit online pharmacy networks. At the same time, Silk Road caters to a priori more technically sophisticated users, and proposes an inventory that far exceeds prescription drugs. As such, it seems to occupy a market niche different from that of traditional illicit online pharmacies. Interestingly, online pharmacies have recently seen their payment systems being targeted [21]. Whether or not this will lead them to move to alternative payment instruments such as Bitcoin, or will redirect their potential customers to marketplaces like Silk Road remains to be seen.

Our work is also close in spirit to a number of studies in the drug policy realm. We particularly single out the work by Caulkins and Reuter [8] who perform an econometric analysis of (offline) drug markets to discuss their structure. Finally, Molnar et al. [23] noted that a number of techniques used to perform analysis of controlled substance markets could be applied to online crime as well. Our paper shows that there may be a convergence between the two fields.

8 Conclusion

We have performed what we believe to be the first comprehensive measurement analysis of one of the largest anonymous online marketplaces, Silk Road. We performed pilot crawls, and subsequently collected daily measurements for six months (February 3, 2012–July 24, 2012). We analyzed over 24,000 items, and parsed over 180,000 feedback messages. We made anonymized versions of our datasets available in a companion website (<https://arima.cylab.cmu.edu/sr/>). We were able to determine that Silk Road indeed mostly caters drugs (although other items are also available), that it consists of a relatively international community, and that a large fraction of all items do not remain available on the site for very long. We further discovered that the number of active sellers and sales volume are increasing, corresponding, when averaged over our measurement interval to slightly over USD 1.2 million/month for the entire marketplace, which in turn represents around USD 92,000/month in commissions for the Silk Road operators. Informed by

these measurements, we discussed some of the possible intervention policies. Using appropriate procedures (e.g., persistent entry guards), Tor shows good resilience to de-anonymization attacks, and, barring operator error, whether it is possible to obtain conclusive evidence of the exact location of a hidden service such as Silk Road remains an open question. Economic attacks (e.g., artificially creating large fluctuations in Bitcoin value), while probably more effective at impeding commerce on such underground marketplaces, would present significant collateral costs. Ultimately, reducing consumer demand (e.g., through prevention campaigns) is probably the most viable strategy.

9 Acknowledgments

We thank our anonymous reviewers, anonymous posters in Bitcoin and Silk Road forums, Gonzague Gay-Bouchery, Steven Kadlec, Tyler Moore, Runa Sandvik, and Zooko Wilcox-O’Hearn for feedback and discussions on earlier revisions of this manuscript. This research was partially supported by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by the National Science Foundation under ITR award CCF-0424422 (TRUST).

References

- [1] The Armory. <http://ayjkg6ombrsahbx2.onion>. Last accessed May 1, 2012.
- [2] Bitcoin charts. <http://bitcoincharts.com/markets/>. Last accessed November 26, 2012.
- [3] Black market reloaded. <http://5onwnspjvuk7cwvk.onion/>. Last accessed November 26, 2012.
- [4] Mt.Gox – Bitcoin exchange. <https://mtgox.com/>. Last accessed July 30, 2012.
- [5] Silk Road anonymous marketplace. <http://silkroadvb5piz3r.onion>. Last accessed November 26, 2012.
- [6] Silk Road forums. <http://dkn255hz262ypmii.onion>. Last accessed November 26, 2012.
- [7] The General Store. <http://xqz3u5drneuzhaeo.onion/users/generalstore/>. Last accessed May 1, 2012.
- [8] J. Caulkins and P. Reuter. What price data tell us about drug markets. *Journal of Drug Issues*, 28(3):593–612, 1998.
- [9] J. Caulkins, P. Rydell, W. Schwabe, and J. Chiesa. *Mandatory minimum drug sentences: throwing away the key or the taxpayers’ money?* RAND, Santa Monica, CA, 1997.
- [10] A. Chen. The underground website where you can buy any drug imaginable. <http://gawker.com/5805928/the-underground-website-where-you-can-buy-any-drug-imaginable>. Last accessed May 1, 2012.
- [11] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, Aug. 2004.

- [12] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. *IEEE Journal on Selected Areas in Communications*, 24(5):1010–1019, 2006.
- [13] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of 14th ACM Conference on Computer and Communications Security (CCS)*, pages 375–388, Alexandria, VA, October 2007.
- [14] D. Houser and J. Wooders. Reputation in auctions: theory and evidence from eBay. *Journal of Economics & Management Strategy*, 15(2):353369, Summer 2006.
- [15] C. Kanich, K. Levchenko, B. Enright, G. Voelker, and S. Savage. The Heisenbot uncertainty problem: challenges in separating bots from chaff. In *Proceedings of USENIX LEET'08*, San Francisco, CA, Apr. 2008.
- [16] E. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53:457–481, 1958.
- [17] N. Leontiadis, T. Moore, and N. Christin. Measuring and analyzing search-redirection attacks in the illicit online prescription drug trade. In *Proceedings of USENIX Security 2011*, San Francisco, CA, Aug. 2011.
- [18] K. Levchenko, N. Chachra, B. Enright, M. Felegyhazi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, A. Pitsillidis, N. Weaver, V. Paxson, G. Voelker, and S. Savage. Click trajectories: End-to-end analysis of the spam value chain. In *Proceedings of the IEEE Symposium and Security and Privacy*, Oakland, CA, May 2011.
- [19] T. May. Crypto anarchy and virtual communities, Dec. 1994. Available from <http://groups.csail.mit.edu/mac/classes/6.805/articles/crypto/cypherpunks/may-virtual-comm.html>. Last accessed July 30, 2012.
- [20] D. McCoy, K. Bauer, D. Grunwald, T. Kohno, and D. Sicker. Shining light in dark places: Understanding the tor network. In *Proceedings of the 8th Privacy Enhancing Technologies Symposium (PETS 2008)*, Leuven, Belgium, July 2008.
- [21] D. McCoy, H. Dharmdasani, C. Kreibich, G. Voelker, and S. Savage. Priceless: The role of payments in abuse-advertised goods. In *Proceedings of the 19th ACM conference on Computer and communications security*, Raleigh, NC, Oct. 2012.
- [22] D. McCoy, A. Pitsillidis, G. Jordan, N. Weaver, C. Kreibich, B. Krebs, G. Voelker, S. Savage, and K. Levchenko. Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs. In *Proceedings of USENIX Security 2012*, Bellevue, WA, Aug. 2012.
- [23] D. Molnar, S. Egelman, and N. Christin. This is your data on drugs: Lessons computer security can learn from the drug war. In *Proceedings of the New Security Paradigms Workshop*, Concord, NH, Sept. 2010.
- [24] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In *Proceedings of the Second APWG eCrime Researcher's Summit*, Pittsburgh, PA, Oct. 2007.
- [25] T. Moore and B. Edelman. Measuring the perpetrators and funders of typosquatting. In R. Sion, editor, *Financial Cryptography and Data Security (LNCS 6052)*, pages 175–191. Springer-Verlag, 2010.

- [26] M. Motoyama, D. McCoy, K. Levchenko, S. Savage, and G. Voelker. An analysis of underground forums. In *Proceedings of ACM Internet Measurement Conference*, Berlin, Germany, Nov. 2011.
- [27] M. Motoyama, D. McCoy, K. Levchenko, S. Savage, and G. Voelker. Dirty jobs: The role of freelance labor in web service abuse. In *Proceedings of USENIX Security 2011*, San Francisco, CA, Aug. 2011.
- [28] S. Mudambi and D. Schuff. What makes a helpful online review? a study of customer reviews on Amazon.com. *MIS Quarterly*, 34(1):185–200, March 2010.
- [29] S. Murdoch. Hot or not: revealing hidden services by their clock skew. In *Proceedings of the 13th ACM conference on Computer and communications security*, pages 27–36, Alexandria, VA, Nov. 2006.
- [30] S. Nakamoto. Bitcoin: a peer-to-peer electronic cash system, Oct. 2008. Available from <http://bitcoin.org/bitcoin.pdf>.
- [31] L. Øverlier and P. Syverson. Locating hidden servers. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 100–114, Oakland, CA, May 2006.
- [32] F. Reid and M. Harrigan. An analysis of anonymity in the Bitcoin system, May 2012. Pre-print. Available as arXiv:1107.452a4v2 [physics.soc-ph]. <http://arxiv.org/abs/1107.4524>.
- [33] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of eBay’s reputation system. *Advances in Applied Microeconomics*, 11:127–157, 2002.
- [34] X. Roche. HTTrack: Website copier. <http://www.httrack.com>. Last accessed July 30, 2012.
- [35] B. Shebaro, F. Perez-Gonzalez, and J. Crandall. Leaving timing-channel fingerprints in hidden service log files. In *Proceedings of the Digital Forensics Research Workshop*, Portland, OR, 2010.
- [36] Y.-M. Wang, M. Ma, Y. Niu, and H. Chen. Spam double-funnel: connecting web spammers with advertisers. In *Proceedings of the 16th international conference on World Wide Web, WWW ’07*, pages 291–300, Banff, Alberta, Canada, 2007.