

Reputation in the Internet black market: an empirical and theoretical analysis of the Deep Web

ROBERT AUGUSTUS HARDY* AND JULIA R. NORGAARD**

Department of Economics, George Mason University, Fairfax, VA, USA

Abstract. This paper is an analysis of the role reputation plays in the Deep Web using data from the Internet black-market site, The Silk Road. This encrypted online marketplace employed cryptocurrency and functioned over the Tor network. Utilizing a modeling technique, informed by trade auction theory, we investigate the effect of seller reputation. Analysis of the seller's reputation gives us insights into the factors that determine the prices of goods and services in this black marketplace. Data on cannabis listings is parsed from the Silk Road website and covers an 11-month time period, from November 2013 to October 2014. This data demonstrates that reputation acts as a sufficient self-enforcement mechanism to allow transactions. These findings exemplify the robustness of spontaneous order with respect to the Deep Web as an emergent marketplace.

1. Introduction

Modern black markets have in place numerous institutions to facilitate trade and evade law enforcement. Cash makes transactions untraceable, hierarchy delineates roles and responsibilities, and violence encourages participants to abide by norms. The advent of the Internet razes this system; entirely new institutions are required for black market trades in this environment. The increased anonymity lowers the risk of detection by law enforcement in exchange for an increase in the risk of impropriety between buyer and seller. This paper examines the use of seller ratings to facilitate trade through lower transaction and information costs.

Illegal Internet activities are conducted on a portion of the Internet referred to as the Deep Web, and is estimated to be thousands of times larger than the Surface Web, the Internet we use every day.¹ The Deep Web is unregulated, untaxed, and hidden from a typical Internet search. It is a self-contained market place that functions under a set of informal institutions. Using a representative data set mined from The Silk Road, one of the most popular sites on the Deep Web,

*Email: rhardy3@masonlive.gmu.edu

**Email: jnorgaar@masonlive.gmu.edu

¹ For a more detailed analysis of the Deep Web and surface web, see Chandler (2013).

we investigate the operation of these black-market transactions. We observe that the institution of seller reputations creates a stable trading environment among those least expected to deal honestly: criminals.

Black-market activity on the Deep Web is attractive because of the anonymity it provides. Cryptocurrencies such as Bitcoin (BTC) function like cash; they are untraceable. The TOR network anonymizes web traffic. PGP encryption programs mask data within emails sent between users. These three elements form the technological base upon which Deep Web black markets build, allowing exchange at a much lower cost than previously. Before this technology, sellers and buyers in the black market relied heavily on face-to-face interaction and building a reputation through personal encounters. This shift led to a flourishing peer-to-peer underground marketplace expanding on a global scale.

But, anonymous Internet trading incurs an additional cost. Like buyers and sellers on any peer-to-peer Internet site such as eBay or Amazon, buyers and sellers on the Deep Web rarely, if ever, meet in person. This makes transactions particularly risky because there is no recourse for failure. And unlike goods on Surface Web sites, Deep Web users are buying products much more harmful than ordinary consumer purchases. The unique nature of this marketplace makes the accumulated reputation of users critical to its emergence and sustainability. Similar to Avner Greif's work on the Maghribi Traders, these Internet traders have asymmetric information (Greif, 1989). However, unlike the Maghribi Traders, these Internet traders have no legal contract enforceability (Skarbek, 2008). Analyzing this reputation component will enlighten, more fully, how this market place can exist without any ability to seek recourse *ex post* and without any prevalence of contract enforceability (Greif, 1993). We empirically answer the questions; Does investment in reputation allow sellers to charge premium prices, or to simply remain in the market? How does reputation play a role in this marketplace?

The most important institution of the Deep Web is anonymity. Each buyer and seller is known by a unique username; their true identity is secret. Users of the Deep Web, through forums and blogs, create a wealth of information to keep users updated on the happenings of the market (DarkNet Markets, 2014). Images 1–3 show Reddit's Deep Web forum and how the users communicate. They use these 'news outlets' to keep users informed on frauds, scams, and imposters. Deep Web markets take a cut of each transaction to cover their operation costs and to make a profit. Buyers write and read extensive reviews on sellers and their products. Markets allow ratings from 0 to 5 stars, accompanied by a brief note explaining the rating. More extensive reviews are commonly posted on internal forums and Reddit. These jointly create the seller's reputation. Some sellers, to differentiate, offer free samples or extra secure shipping techniques to attract positive reviews.

This paper investigates a market place where feedback mechanisms and reputation are the only things keeping the market functioning, without any

government taxation and regulation (Clay, 1997; Greif, 1989). Deep Web markets are an empirical example of the depth of robustness of spontaneous order. It shows that the principles of an unfettered market rooted in reputation and accountability can be applied to an extremely vast array of goods and services. We are fundamentally analyzing how individuals interact with each other and without government (Leeson, 2010; Powell and Stringham, 2009). In Section 2, we delve into the factors that differentiate the Deep Web from other online marketplaces. Section 3 explores how reputation provides a market mechanism to keep buyers and sellers accountable and honest. We outline and describe our theoretical model in Section 4, analyzing how reputation functions in the market. Our empirical method of analysis is laid out in Section 5. Section 6 includes a description of our data, our collection procedure, and detailed definitions of all our variables. Section 7 reviews our results and estimates the buyers' and sellers' discount factors. Our concluding remarks about the implications and impact of our findings are enclosed in Section 8.

2. An overview of the Deep Web: what differentiates it from other online marketplaces

The currency used to make transactions in the Deep Web is BTC. A BTC is a solution to a mathematical equation, and a pseudo-anonymous cryptocurrency (Grinberg, 2011). They are stored in virtual wallets and are exchanged through anonymous virtual transactions with low transaction fees (Briere *et al.*, 2013). To 'mine' BTCs, 'miners' use computing power to solve mathematical problems to which there are a fixed number of solutions. Because of the fixed nature of the number of possible solutions (BTCs), this cryptocurrency, by design, cannot be inflated. Therefore, this alternative currency is free from central bank policies or intervention (ECB, 2012).

According to Christin (2012), 'BTC 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'. BTCs are used for Deep Web transactions because they are anonymous, like cash, and can be transacted electronically. A BTC wallet functions like a physical wallet with cash: once you transfer BTCs from one wallet to another, it is untraceable and permanent.

The Deep Web exists on Tor, a computer networking system that allows for anonymous communication and transactions. The communications sent on Tor are encrypted and then sent through numerous network servers and nodes. When users communicate through the Surface Web, their messages are unencrypted and travel directly from sender to receiver. Messages are 'bounced' between nodes in the Tor network, making them virtually untraceable. The random path the message takes, coupled with its encryption while traveling through the network secures the anonymity of the users and security of the content. This message 'bouncing' cause the Tor network to be much slower than Surface Web networks.

The identity of the sender and receiver of a message over Tor is hidden unless the user explicitly wishes to reveal their identity (Onion Routing, 2014). Because BTC is an anonymous cryptocurrency, it is used as a medium of exchange on Tor.

The barriers to entry into the deep web are very high. The use and knowledge of BTC takes some degree of computer sophistication. However, BTCs are becoming increasingly popular and information about how to obtain and use them is readily available. The use and implementation of Tor, on the other hand, suffers from a very large knowledge problem. Jeffrey Tucker, founder of Liberty.me, describes the skills it takes to feasibly and securely make transactions over the Deep Web; ‘you have to be a sophisticated person to get into commercial buying and selling on the Silk Road’ (Tucker, 2014). Users thus turn to this type of market place because it provides them with goods and services more cheaply, more safely, or of a higher quality than their local black market would allow them to access. According to Tucker, ‘People have an intensity of demand to overcome technical barriers’ because there are no online tutorials and much of what goes on in Tor is illegal. There are also very high risk factors when it comes to anonymizing oneself, detection of one’s identity could result in stolen goods, personal safety issues, or imprisonment.

A primary difference between traditional online sites, such as eBay, and the Silk Road is escrow implementation. Standard escrow requires the ability to undo a transaction. Fraudulent items are returned to the seller, and then the escrow service refunds the buyer. Hu *et al.* preface their model on the assumption that ‘in the case of fraud, [escrow] users lose only the service fee’ (Hu *et al.*, 2004). Silk Road purchases cannot be undone; drug dealers do not provide return addresses. An escrow service cannot exist which simultaneously satisfies buyer and seller.

The Deep Web is an untaxed and unregulated marketplace; it exists as a completely unfettered free-market. This marketplace functions much like the historical Law Merchant market did in medieval Europe and the medieval Maghribi Traders in the Mediterranean (Greif, 1989; 2012). According to Benson (1989)

the rules of property and contract necessary for a market economy, which most economists and legal scholars feel must be ‘imposed’, have evolved without the design of any absolute authority. Commerce and commercial law have developed conterminously, without the aid of and often despite the interferences of the coercive power of nation-states because there is a mechanism in place (1989: 644–645).

With respect to the Silk Road, the ‘internal policing’ mechanism that Benson refers to is the reputation of sellers and buyers.

Because the users in this marketplace cannot seek legal recourse for their illegal transactions, they must police themselves (Milgrom *et al.*, 1990). The Deep Web Culture promotes transparency with respect to the quality of the

goods and services as well as honesty amongst buyers and sellers. Users have created checks and balances on each other to feel confident and safe on the Deep Web. Just like historical pirates (Leeson, 2007), buyers use checks and balances to constrain seller predation. In the absence of a central coercive force for recursive action, users must rely on each other for feedback and information. The security and reliability of this network is what keeps users confident in the marketplace because they provide internal checks on each other. Many forums contain information about people who are masquerading as prominent sellers, or users that are committing fraud. This emergent order is no surprise. According to Leeson (2010) and Mueller (1988), organization and structural norms emerge without the use of a central planner in the marketplace and these norms are effective at keeping users in the marketplace safe and satisfied with their services and products.

This marketplace has allowed for anonymous peer-to-peer engagement with only the Silk Road and other hosting sites to facilitate the exchange and take a small fee. Resnick and Zeckhauser (2001) stress that when there is repeated play among individuals in a marketplace it reduces the likelihood of dishonest people continuously dealing in the market and reduces moral hazard. This type of transaction has revolutionized the illegal goods and services trade because it has made it more convenient, accessible, and has allowed users to access a larger variety of the good of their choice. This online network has enabled local sellers of illegal goods to expand to a global setting, and increasing worldwide price and quality competition. In terms of the global drug market, the Silk Road is a small fraction. Kilmer and Pacula (2009) estimate a 2003 trade volume of \$142 billion. Court documents used in the trial of Silk Road Founder Ross Ulbricht (U.S. vs. Ross Ulbricht, 2013) allege the original Silk Road grossed approximately \$214 million during its two years of operation.

The latest estimate of marijuana street prices comes from the Office of National Drug Control Policy (ONDCP). The ONDCP uses data from the National Survey on Drug Use and Health (NSDUH) to predict price per gram on the street. They find relatively stable prices, though the 2010 estimate of \$7.11/gram has declined from 2004's \$7.50/gram. They also extend the analysis to Fries *et al.* data set, estimating a 2010 street price of \$10.70/gram (Fries *et al.*, 2008). Our data set's median 2014 price of \$13.61/gram is consistent with a number of theories: higher quality marijuana, an Internet premium, or price inflation over the past four years. Morgan (2012) finds that revenue lost to online fraud is falling, estimated to 0.9% in 2011.

3. Reputation as a mechanism for market accountability

Because of the nature of the goods sold in the Deep Web, on the Silk Road in particular, sellers are anonymous to buyers and buyers are anonymous to sellers. Before a first transaction, they have no personal knowledge of another's

personality and no formal enforcement mechanism if a transaction goes awry. The characteristics of this particular marketplace pose risks to the traders involved. The buyer could refuse to pay the seller after their items have been received, or, if the buyer pays first, the seller could fail to send the purchased items because they received the payment upfront. There is no way to recoup lost BTCs or products once the transaction is finalized. This marketplace exists due to the importance of a bilateral reputation mechanism that instills confidence in the traders and facilitates repeated transactions (Greif, 2012).

Leeson (2005) makes the point that users 'need to establish *ex ante* whether or not the outsiders they would like to trade with are 'cheaters' or 'cooperators'. In other words, they need a means of screening outsiders' (Leeson, 2005: 79). Collecting as much data as possible on the other party is necessary to making a smart and calculated transaction. Initially, buyers and sellers are dependent upon previous users' feedback for information on the legitimacy of their potential trade. Recognizing this potential risk, traders utilize forums such as Reddit and the Silk Road itself for feedback, bringing attention to fraudulent behavior and informing traders of transaction malfeasance.

The codification of buyer and seller feedback makes up each party's user profile (Houser and Wooders, 2006). A user's feedback profile in this marketplace is made up of the comments and ratings left on the Silk Road site as well as other feedback forums. This feedback is both comments and a number rating. The collection of this user feedback on other users makes up the reputation of the trader in the marketplace. Due to the anonymity aspects of The Silk Road, buyer information is not formally posted like seller information and feedback is on the site. Unlike Surface Web marketplaces, if a buyer leaves a comment and/or rating, an individual identifier is not attached to their message. The reason for this is to protect buyer anonymity. The only information that we can glean about the buyer in particular is that they did in fact make a purchase; buyers cannot leave feedback on a product they did not buy.

Potential buyers utilize this feedback about sellers. They can read comments about previous buyer's experiences, whether or not the buyer received the items, and view the seller's 30-day and 60-day and overall rating score. This score is an average of past reviews and it is out of five possible points. Sellers, however, do not have access to this information about potential buyers. Repeated trade will reveal buyer reputation, but the first is made with little information. The promise of future trade can incentivize honest behavior from the beginning; sellers can cease trade with dishonest buyers.

Discovery of a dishonest buyer can have positive externalities for other sellers. But, sellers' outlets for relaying the information that they have learned from buyers are limited. Because buyers do not have publically available profiles, the seller must seek alternative forms of feedback. They can leave feedback on the internal Silk Road forums or various forums on the Surface Web, but cannot add to a collected reputation buyer profile because they do not exist.

Gambetta (2010), identifies that criminals need both a costly signal of the trader's credentials and a costless arbitrary group signal in order for this type of market place to run smoothly underground. Leeson (2005) further breaks down the components necessary for a successful reputation signal in general. He states that they must be easily observable and that they also must be costly for cheaters to signal a stellar reputation and fairly inexpensive for honest users to signal that they are authentic.

Applying these characteristics to the Silk Road marketplace, the seller feedback mechanisms of readily observable ratings, comments, and thus reputation fit these criteria and send a signal that the seller is honest or dishonest. It would be difficult for a repeatedly dishonest seller to trick its buyers to leave positive reviews and ratings even though the products and services were a sham. On the other hand, if an honest seller provides their customers with quality products in a timely manner, it will be relatively easy to receive truthful positive reviews about the seller's quality performance. This dovetails very nicely with what we know about the Silk Road community from studying Silk Road forums: the community is very active at giving feedback. These criteria, easily observable signaling and costly signaling for cheaters, do not necessarily apply to the buyers in this marketplace. This failure of buyer feedback to meet the strong signal criteria proposes that buyer signals could contain a great deal of noise and potential for misread signals. For the purposes of this paper, we will analyze the impact of seller's reputation as a signal.

However imperfect these feedback mechanisms may be, they provide users information on reputation. Reputation is crucial in this market because it acts as a signal to other users that they are honest and credible individuals. This signal works to differentiate between honest and dishonest users to ensure that honest users are not driven out of the marketplace by dishonest users that are not properly identified. Leeson (2005) emphasizes that the traders' identities work to reduce social distance in the marketplace. Deep Web traders do not have an identity in the traditional sense; however, they foster an identity through their online reputation. The Leeson model makes it clear that people cheat because they have higher discount rates than their cooperators. Their gains from future exchange are more heavily discounted, thus they invest less because it is more costly for them.

Our analysis in this paper estimates the discount factors of all users. An essential component to the reputation system is that, if reputation does allow sellers to charge their customers a premium, it behooves the sellers to increase their reputation so as to be able to collect premium profits. Therefore, the existence of the reputation system itself acts to ensure honesty with each transaction. This particular phenomenon is what this paper analyzes in great detail, whether or not an increase in reputation empirically and statistically significantly allows sellers to actually charge premium prices. We analyze if

favorable reputation allows sellers to capitalize on their positive feedback and signal to buyers that their items are of high quality like Shapiro (1983) found.

4. Theoretical and empirical model

We seek to accomplish two goals. First, we disentangle the role of reputation on the Silk Road. Does investment in reputation allow sellers to charge premium prices, or to simply remain in the market? Houser and Wooders (2006) posit a market with *honest* and *dishonest* sellers. Reputation serves as a signal that a seller is honest. Thus, a buyer's utility, as well as willingness to pay, increases with increased seller reputation. This reputation enables sellers to earn a premium. Houser and Wooders find evidence supporting this theory using their data on eBay auctions. Other papers find similar results. Klein and Leffler (1981) examines the use of higher prices to ensure contractual performance, Shapiro (1983) as well as Allen (1984) examine prices above marginal cost to forestall quality cutting, and McDonald and Slawson (2002) examine returns to reputation in electronic auction markets.

Alternate theoretical models can be constructed such that reputation does not convey a premium. Rather, in equilibrium all sellers are *honest*. One could suppose that above a certain threshold a seller is considered *honest* by buyers and remains in the market, below that threshold the seller is considered *dishonest* and exits the market or creates a new identity. Melnik and Alm (2002) find some support for this theory. They show a positive relationship between reputation and price, but the predicted effect is quite small.

Second, what assumptions about time discounting must be made to sustain the market? Using the estimates of return to honesty, we can determine a lower bound to a seller's time preference. A seller with a higher time preference would be less patient for payment and would prefer the buyer to pay for the goods before they had received them. This is called finalizing early, and occurs when buyers transfer their payment in BTCs to the seller before the product is received. If the seller has a relatively low time preference, they would not necessarily request a finalize early payment. Not much work has been done on estimating this variable, though it has vast implications for the functioning of a market dependent upon repeated trade and weak punishments. Finalizing early was prevalent on the Silk Road prior to its shutdown. As markets have evolved, multi-signature escrow has become the norm.

We begin with a simple model, discussing the interaction between a single seller and buyer. We later expand this simple model's insight to a broader model. Choice nodes exist for buyers and sellers, each dependent upon the expectation of actions at the subsequent nodes. First, we analyze the nodes chronologically and then, by backwards induction, create a theory to predict market action and the general qualities of equilibrium.

We posit a good j sold by seller j to buyer i . The item is listed at a price p , and has a value of V_i to the buyer. Both buyer and seller have initial, publicly known, reputations r^b and r^s . Reputation serves as a proxy for the probability that the individual will act honestly. The model incorporates a signal extraction problem: honest behavior can be perceived as dishonest. A package may fail to arrive because it was intercepted by law enforcement (honest), or because it was never shipped (dishonest). We create a variable, r^c , representing signal clarity. r^c takes a value from 0 to 1; 1 implies perfect signal transmission, and 0 complete signal failure. There is a probability, $(1-r^c)$, that an honest signal is received as dishonest. Reputations are therefore imperfectly updated. Production for a single unit of good j costs c , which includes production costs as well as shipping costs in most cases (many sellers offer free shipping). Finally, actors discount future periods by β_i and β_j . Each of these discount factors depends on the buyer and seller's time preferences.

The seller takes the first step, creating a listing. The seller sets all aspects of this listing: product, price, and method of payment. Product description and price have an unbounded set of possibilities, and equilibrium occurs within the intersection of this possibility set and the buyer's demand set. Assume that buyer and seller interact within this intersection. Action outside of this intersection is uninteresting; no trade occurs.

The meaningful choice we are left with at this node is method of payment. The seller chooses what occurs first: buyer payment or seller shipment. If the seller requests that the buyer finalize early, the buyer pays for the product before shipment. We analyze buyer pays first.

The buyer's decision in the case of 'buyer pays first' is simple. Do I value the item above the cost? Similar to Houser's treatment of reputation and value, any purchase must satisfy the equation

$$p \leq r^s V_i \tag{1}$$

The expected benefit to the buyer must be greater than the price of the item. If this equation holds, the buyer will make the purchase. Otherwise, no transaction occurs. We assume, in equilibrium, that the seller will raise price until the previous equation is binding, that is

$$p = r^s V_i \tag{2}$$

The seller now faces the decision to be honest or cheat. If the seller is honest, item j is shipped and the seller's reputation increases. If the seller cheats, item j is not shipped and the seller's reputation decreases. Because price is a function of seller reputation, honest sellers can charge premiums. The price that an honest seller can charge is p_h and the price that a cheating seller can charge is p_c . Cheating once followed by honesty results in a payoff of $p + \sum_{t=1}^{\infty} \beta_j^t (p_c - c)$. The payoff to honesty is thus $(p - c) + r^c \sum_{t=1}^{\infty} \beta_j^t (p_h - c) + (1 - r^c) \sum_{t=1}^{\infty} \beta_j^t (p_c - c)$.

The actions of the seller alter his reputation, such that it increases with honest action and falls with dishonest behavior. The reputation of an honest seller is thus r_h^s and the reputation for a dishonest seller is r_c^s . Put mathematically, $r_c^s < r^s < r_h^s$. Recall that in equilibrium $p = r^s V_i$.

Thus,

$$\sum_{t=1}^{\infty} \beta_j^t (p_c - c) < \sum_{t=1}^{\infty} \beta_j^t (p_h - c).. \quad (3)$$

$$R_h = \sum_{t=1}^{\infty} \beta_j^t (p_h - c) \quad (4)$$

$$R_c = \sum_{t=1}^{\infty} \beta_j^t (p_c - c) \quad (5)$$

Equation (3) demonstrates that the expected future revenue stream from honesty is strictly larger than that from cheating. This is because honesty raises a seller's reputation, allowing seller j to charge premium prices indefinitely. These revenue streams are simplified to R_h and R_c in equations (4) and (5) respectively. We can now analyze the conditions under which a seller will remain honest.

$$\begin{aligned} (p - c) + r^c R_h + (1 - r^c) R_c &\geq p + R_c \\ r^c (R_h - R_c) &\geq c \end{aligned} \quad (6)$$

A seller remains honest when the expected cost of cheating is greater than or equal to the cost of production. Honesty is increasing with both signal clarity and future returns to high reputation. It is decreasing with future returns to cheating and cost.

A buyer has no way of knowing that equation (6) holds, instead inferring the seller's honesty through r^s . The buyer's decision being unaffected by equation (6) suggests that equilibrium is reached. Attention should be paid to three facets of this equilibrium. As previously mentioned, the seller will raise price to meet the buyer's valuation, $p = r^s V_i$. Increases in the return to honesty, $(R_h - R_c)$, will make the equilibrium more stable. The signal of an honest seller will be stronger if sellers are incentivized to be honest because the signal clarity is important for trade reoccurrence.

Now, the case of 'seller ships first.' The buyer still purchases when $p \leq r^s V_i$, but the seller now assumes all risk; r^s is effectively 1. If the expected product is not delivered, the buyer will withhold payment at no loss.

The seller's choice is to engage in the transaction, or decline. When the expected revenue exceeds the expected costs, the seller engages in the transaction

$$r^b V_j \geq r^b p \geq c \quad (7)$$

A seller remains in the market while equation (7) holds. Unlike r^s , r^b is not known. Buyers do not have public reputations, thus r^b is the average expected buyer reputation. After receipt of good j , buyers choose to cheat or be honest. Cheating buyers withhold payment, and receive a payoff

$$V_i + (1 - r^c) \sum_{t=1}^{\infty} \beta_i^t (V_i - p) \quad (8)$$

Honest buyers receive payoff

$$(V_i - p) + \sum_{t=1}^{\infty} \beta_i^t (V_i - p) \quad (9)$$

When the seller does not receive payment, either the buyer is cheating or the good was intercepted. If seller j believes the buyer is honest, trade may occur again. Combing both equations gives conditions of buyer honesty.

$$r^c \sum_{t=1}^{\infty} \beta_i^t (V_i - p) \geq p \quad (10)$$

This inequality states that when the future benefits from trade; discounted by signal clarity, exceed price the buyer will behave honestly. It is important to note that lowering price unambiguously makes equation (10) more likely to hold. This incentivizes the seller to lower price until equation (7) binds

$$r^b p = c \quad (11)$$

A few qualities of this equilibrium emerge. Honest buyers enjoy a surplus of $V_i - p$, because a surplus of 0 would cause equation (10) to not hold. This surplus allows the market to function, by rewarding honest buyers at the expense of cheaters. Price is a function of production costs as well as buyer reputation. Market durability is increasing in r^c and β_i . V_i can have positive effects, but it depends upon the marginal cost increase of an increase in value.

The addition of more buyers and sellers strengthens the market. The equalities previously derived now apply to the marginal buyer and seller. Different V_i 's, β_i 's, and β_j 's may allow some to benefit more than others in equilibrium. In the case of buyer 1st, equation (2) will still hold. The increase in market size will likely drive equation (6) to the binding point. Seller entry will put downward pressure on the returns to reputation, until entry ceases at the binding point.

$$r^c (R_h - R_c) = c \quad (12)$$

Plugging equations (4) and (5) into (12) generates an enlightening equation for seller patience

$$r^c \left[\sum_{t=1}^{\infty} \beta_j^t (p_h - c) - \sum_{t=1}^{\infty} \beta_j^t (p_c - c) \right] = c$$

$$\frac{r^c \beta_j}{1 - \beta_j} (p_h - p_c) = c \quad (13)$$

$$\frac{\beta_j}{1 - \beta_j} = \frac{1}{r^c} \frac{c}{p_h - p_c}$$

The marginal seller in equilibrium must discount the future such that (13) holds.

Similar price pressure will occur in the case of seller 1st. The price equation, (11), still holds. Taking equation (10) further calculates an equality similar to (13)

$$\frac{\beta_i}{1 - \beta_i} = \frac{p}{V_i - p} \frac{1}{r^c r^e} \quad (14)$$

The marginal buyer has a β or V_i low enough to make equation (10) binding.

5. Empirical method

Recall that in ‘buyer pays first’ equilibrium, the price of j is given by equation (2).

Thus, by taking logs of both side

$$\ln(p) = \ln(r^s) + \ln(V_i) \quad (15)$$

The log of price is a function of the observable seller characteristics, r^s , and the observable item value, V_i . Because our data contains sellers of multiple products, we expect heteroskedastic errors correlated by seller. This equation can thus be estimated using generalized least squares and standard regularity conditions. We follow the basic estimation method used in Houser and Wooders (2006).

In addition, equation (13) can be estimated. We do not have sufficient data to attempt equation (14), because buyer’s personal values cannot be directly observed. To empirically estimate equation (13), we first make a slight transformation.

$$\frac{\beta_j}{1 - \beta_j} = \frac{1}{r^c} \frac{c/p}{[(p_h/p) - (p_c/p)]} \quad (16)$$

Our previous regression will estimate for us $\frac{p_h}{p}$, the percent increase in price given one more positive review. Assuming linearity in returns to reputation, the percent decrease in price given a negative review will equal the increase given a

positive review. Thus,

$$\frac{\beta_j}{1 - \beta_j} = \frac{1}{r^c} \frac{c/p}{(2^* p_h/p)} \quad (17)$$

In a market with mixed payment methods, such as our data set, we assume that price fulfills a combination of equations (1) and (7), price is less than or equal to $r^c V_i$ and greater than or equal to c/r^b . This means our estimates will be imperfect, but can be checked later for robustness. In theory, competition will minimize this gap. Therefore,

$$\frac{\beta_j}{1 - \beta_j} \approx \frac{1}{r^c} \frac{r^b}{(2^* p_h/p)} \quad (18)$$

We later estimate r^b and r^c .

6. Data from the silk road

We use sales data on 119 cannabis listings from 41 sellers, for a total of 9,604 sales. Transaction volume ranges from a single sale to 688. Though they can be used in a variety of ways, cannabis products have a single purpose. We assume that different strains are highly competitive, and similarly that if a well-recognized strain exists it does not command a premium. Our data set is parsed from the Silk Road website and covers an 11-month time period, from the opening of Silk Road II in November 2013 to our collection date in October 2014. We chose cannabis sales exclusively because it is one of the biggest portions of the market with a lot of differentiation of product type and strain. We also want to look into one type of market, presuming the marketplaces for other products are differentiated.

It is important for our empirics that we use data on sales of marijuana for personal use. Our theoretical model requires that item value assessments must use private values, not common values. In Virginia, the cutoff for misdemeanor possession charges is less than 1/2 oz, or approximately 14 grams. Anything above 1/2 oz is considered intent to sell, and carries a felony charge. By contrast, Florida draws the line at 20 grams. Our data has a fairly natural break at 15 grams, so we will perform empirics on weights of 15 grams and below, ensuring that we ignore listings meant for resale.

Data collection procedure

Data was collected from the Silk Road (silkroad6ownowfk.onion) using a web crawler called HTTrack. HTTrack utilized the Tor network to download the web pages and structure of the Silk Road over the course of four days in early October. This was a slow process; the Tor network limited the download to around 4 KB/sec. To economize on bandwidth, the download ignored all images and only downloaded the text of web pages. Appendix A includes sample web

page images. The downloaded webpage data was then used to create a local mirror of the site.

Data was then parsed from the site into an Excel file using a custom parser. This created three unique data sets: seller data, feedback data, and item data. Our focus for this paper is the item data. Our parser gathered data on the listing name, price, aggregate item feedback values, aggregate seller feedback values, free shipping, number of sales, days sold for, and weight.

We will now outline some of the major difficulties with the data. The crawler is imperfect, and known to make mistakes. Of the over 30,000 files downloaded, HTTrack reported approximately 300 errors. Given the extended download period, nature of the connection and size of the download, these errors were expected. The key is that they are random. Errors typically occurred singly, at a rate of a few an hour. Also, few errors completely eliminated information on a seller or listing. Rather, they removed a page of feedback details for a user. Finally, errors did not change any data points; it instead makes them unreadable.

Price data also presented complications. BTC prices can fluctuate dramatically. To alleviate this, the Silk Road appears to pin listing prices to some more stable currency. The mechanism is not public knowledge, but we theorize that seller's indicate a price in USD that is then converted to BTC. This value is then periodically updated as the conversion rate changes. This creates problems when downloading over multiple days, as relative prices change due to adjustments in BTC exchange rates rather than value. To correct for this, the parser converted BTC prices to USD based upon the date of download and exchange data from the Coindesk.

Weight is our restrictive variable. The Silk Road provides no universal way to list the products weight, which creates complications for the parser. Some listings show weight in the title, others somewhere in the description. Different countries deal with decimals differently; one may list 3.14 grams, another 3,14 grams. Finally, 'grams' can be abbreviated as g, gr, or omitted entirely in the case of 'sample packs'. The parser was created to grab as many weights as possible, but could not grab them all. For future work, we will likely enter more data by hand.

The key difference in our model between 'buyer pays first' and 'seller pays first' is the price function. In 'buyer pays first', price is determined by the reputation of the seller and the value of the good. In 'seller pays first', it is a function of production cost and the reciprocal of buyer reputation. The Silk Road is a mixture of both systems.

Silk Road uses two payment methods: finalize early and an escrow system. Finalize early closely approximates our Buyer 1st model previously outlined. A seller requires that a buyer complete payment before the item ships. Escrow is similar to Seller 1st, but imperfect. A buyer indicates willingness to purchase by sending funds to the Silk Road's escrow account. When the item arrives, the buyer is then expected to release the funds to the seller. If the item does not

arrive, the Silk Road fully or partially refunds the buyer. Typically, this favors the buyer, who receives a full refund. In terms of our model, this system works to increase r^b and r^e by raising the cost of cheating.

It is difficult to determine what payment method a seller offers. Many shift over time, depending on the item, or perhaps offer both (with bonuses to finalizing early). Our data covers a unique period in the Silk Road's history; the centralized escrow account was hacked on February 13th of 2014. This pushed sales to mandatory finalizing early for approximately two months before allowing individual choice again.

Variables used in empirical analysis

Our dependent variable, *lnPricebyWgt*, is the log of an item's list price divided by its weight in grams. As previously mentioned, the price is converted to USD from BTC using Coindesk price data. Weight is converted from its list unit to grams. This creates a value in USD/gram. Listings with no feedback, either because no sales have been made or feedback posted (we cannot differentiate between the two), are ignored.

Our seller reputation variables are *VendorRatingOverall*, *VendorRepFall*, *ItemRatingOverall* and *ItemRatingSum100*. *VendorRatingOverall* aggregates seller feedback. Sales are concluded when a buyer leaves a feedback rating, from 0 to 5 stars. This feedback is averaged over the life of the listing and reported on an item and seller's listing page. *VendorRepFall* is a dummy variable capturing dynamic effects of reputation. Buyers observe three measures of reputation: overall, 60 day, and 30 day. *VendorRepFall* takes a value of '1' if 30-day reputation is less than overall reputation. We include *ItemRatingOverall* and *ItemRatingSum100* as seller reputation variables though they appear to be value related. First, Item Rating and Vendor Rating are highly correlated; Vendor Rating is an aggregate of a sellers' Item Ratings. In addition, feedbacks do not typically include value assessments. They are a bimodal distribution: 5 if the good is received, 0 if it is not. Thus, ratings better reflect a seller's honesty than the goods quality. *ItemRatingSum100* is the sum of feedbacks for the item. This is calculated by multiplying overall feedback by number of feedbacks, then dividing by 100 to make the numbers more manageable. This coefficient will be used to estimate (18). It is worth noting that we do not include number of feedbacks in our regressions due to its strong correlation with *ItemRatingSum100*. Finally, *ItemRepFall* is a dummy variable similar to *VendorRepFall* for the item rating.

Our item value variables are *WgtGrams*, *FreeShip*, *NumFeedbacks*, and *Advert*. *WgtGrams* is a variable, measured in grams, which controls for the weight of the product. This controls for quantity discounts *FreeShip* is a dummy variable, which takes a value of '1' if the seller offers free shipping. It is difficult to disentangle the various shipping options offered by seller. This variable captures the added value of free shipping. *NumFeedbacks* is the number of feedbacks

Table 1. Summary stats

	Variable	Obs	Mean	Std. Dev.	Min	Max
Dependent variable	PricebyWgt	119	14.88	5.56	4.06	32.04
	PriceUSD	119	58.39	51.48	4.30	360.76
Seller reputation variables	VendorRatingOverall	119	4.84	0.20	4	5
	VendorRepFall	119	0.57	0.50	0	1
	ItemRatingOverall	119	4.86	0.21	4	5
	ItemRatingSum100	119	3.94	0.49	0.04	33.64
Item value variables	Advert	119	0.15	0.36	0	1
	FreeShip	119	0.66	0.48	0	1
	NumFeedbacks	119	80.71	108.59	1	688
	WgtGrams	119	4.70	4.47	0.5	15
	ItemRepFall	119	0.40	0.49	0	1

received on this listing. Higher sales numbers may signal to potential buyers that the product is as advertised. *Advert* is a dummy that takes a value of one if the following words appear in the listing name: Premium, AAA, High Grade, Top Quality, or Strong. Marijuana's varying levels of quality are difficult to measure. This dummy weakly controls for quality differences among listings. These only hold analytical strength if we assume both that buyers can tell the difference between high- and low-quality product, and sellers communicate quality through these words.

7. Results

GLS

The results of our GLS regressions estimating (15) are presented in Table 2. We report seven regressions, exploring different measures of seller reputation and controls. The regressions provide evidence supporting the hypothesis that reputation provides a premium.

Our estimates of *ItemRatingOverall* are significant across all the regressions it is included in. The magnitude of the coefficient overshadows our other controls. In contrast, *VendorRatingOverall* is only significant in regression (4). It is insignificant and has the opposite of expected sign in the presence of *ItemRatingOverall*, suggesting that most reputation information is contained within the item's rating rather than the vendor's rating. *ItemRatingSum100* is insignificant in all regressions, though the addition of more controls raises its Z score. It maintains the expected positive sign throughout.

VendorRepFall and *ItemRepFall* return expected results. Given the item value controls, they are significant and negative. Despite the insignificance of *VendorRatingOverall*, changes in the rating are related to changes in the price.

Table 2. GLS regressions of effect of reputation on log price per gram

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>VendorRatingOverall</i>	-0.098 (0.50)	-0.109 (0.55)	-0.130 (0.64)	0.477 **(4.62)			
<i>ItemRatingOverall</i>	0.693 **(4.14)	0.694 **(4.09)	0.728 **(4.2)		0.547 **(8.06)		0.546 **(8.30)
<i>ItemRatingSum100</i>		.0027 (0.63)	.0037 (0.73)			.0056 (1.48)	.0049 (1.68)
<i>VendorRepFall</i>			-0.000 (0.00)	-0.079 *(2.33)	-0.062 *(2.31)	-0.087 **(2.61)	-0.066 *(2.44)
<i>ItemRepFall</i>			-0.032 (0.61)	-0.080 **(3.08)	-0.067 **(3.18)	-0.088 **(3.38)	-0.082 **(3.61)
<i>FreeShip</i>				-0.145 **(4.45)	-0.131 **(4.47)	-0.110 **(2.86)	-0.130 **(4.53)
<i>WgtGrams</i>				-0.044 **(14.48)	-0.044 **(15.76)	-0.044 **(14.64)	-0.042 **(15.24)
<i>Advert</i>				-0.193 **(5.86)	-0.150 **(3.52)	-0.214 **(6.61)	-0.148 **(3.59)
Constant	-0.264 (0.74)	-0.225 (0.61)	-0.278 (0.65)	0.733 (1.41)	0.357 (1.05)	3.015 **(60.54)	0.343 (1.03)
Observations	119	119	119	119	119	119	119
Number of VENDOR	41	41	41	41	41	41	41

Absolute value of z statistics in parentheses.

*Significant at 5%; **Significant at 1%.

Changes in seller rating may be more important than the actual rating because of the inherent risks of these transactions. Sellers face forces largely outside of their control (law enforcement) and can expect a percentage of shipments to be intercepted. A stable, less than perfect, seller rating reflects this. When buyers observe a fall in a vendor's rating, there is thus confusion about the source: seller behavior or law enforcement action.

Estimates of *WgtGrams* are as expected, negative and significant. This indicates a bulk discount of sorts; the more grams purchased the lower the price per gram. *Advert* and *FreeShip* are both significant and negative. We theorize that this is because lower quality (and thus price) goods compete on more margins than high-quality specialized goods. Free shipping and quality descriptions are thus more likely for a low quality good to entice buyers from competitors' products.

Estimating beta

Now we use the estimated coefficient of *ItemRatingSum100* to estimate β_j . The coefficients are not significant, but they are consistently positive across regressions, which is what we predicted, and approximately the same

Table 3. β , given a p_h/p of 0.00245

$r^c \backslash r^b$	0.972	0.968	0.9	0.75	0.5
0.99	0.9,995	0.9,995	0.9,995	0.9,994	0.9,990
0.95	0.9,995	0.9,995	0.9,995	0.9,994	0.9,991
0.9	0.9,995	0.9,995	0.9,995	0.9,994	0.9,991
0.75	0.9,996	0.9,996	0.9,996	0.9,995	0.9,993
0.5	0.9,997	0.9,997	0.9,997	0.9,997	0.9,995

magnitudes. We will show that the actual estimate matters little at this point.

$$\frac{\beta_j}{1 - \beta_j} \approx \frac{1}{r^c} \frac{r^b}{(2^* p_h/p)} \quad (19)$$

We begin using the *ItemRatingSum100* coefficient of 0.0049; regression (7) estimated the highest Z score for the variable. We first divide by 100, to determine the estimated return to a single-star review. A positive review is typically 5-stars, so we then multiply this number by five to determine the estimated percent increase in price given a five-star review. This gives us $p_h/p = 0.000245$.

Do these results suggest that we should we expect the quality of buyers to significantly differ from that of sellers? Perhaps, sellers have stronger incentives to be honest; they must make a larger investment in reputation. But because we cannot directly measure r^b given our current data set, we must make some assumptions. We thus calculate β_j assuming that r^b is the average *ItemRatingOverall* (0.972), the average *VendorRatingOverall* (0.968), and some lesser values. Though we cannot explicitly say that these ratings approximate buyer reputation, we can weakly say that they are an upper bound.

Similarly, we do not currently have good estimates of r^c . We expect signal clarity to vary greatly depending on location, package size, even the time of year. We thus see how various values of r^c affect β . An interception rate of 1% gives a value of 0.99, 5% gives a value of 0.95, etc. For example, a 5% interception rate means that for every hundred illegal packages moving passing through the postal service, five are confiscated. Further work needs to be done to estimate the variable.

As Table 3 shows, the low return to honesty appears to overshadow any considerations of r^c and r^b . This incredibly high β is intuitive given the frequency of transactions. Given this is the time preference between sales, which occur relatively frequently, we should expect the seller to not discount very much. A retailer does not discount sales that will occur in the evening relative to sales in the morning. Sales are typically thought of in terms of days or weeks, not individually. As a thought exercise, the average seller has been active for 10 months, and the average listing has received approximately 80 feedbacks. This reduces to an average of eight feedbacks per month, or two feedbacks per week,

or one feedback every 3.5 days. Approximately 104 3.5-day sets occur every year. By taking our high (0.9,997) and low (0.9,990) estimates of β to the 104th power, we can estimate an annual β range. This works out to be $0.9012 \leq \beta \leq 0.9693$, and can be converted to an interest rate so that $0.0317 \leq r \leq 0.1096$. More robust analysis of these estimates is required.

8. Conclusion

We note some possible objections to our model. This model ignores the possibility of utility of action. Utility of actions means that an individual may receive utility from the action of honest behavior. This can be added to the model by creating a constant value to one of the payoffs. A positive constant added to the payoff for honesty or payoff for cheating could simply model utility of honesty or cheating respectively. This possibility makes any estimates of β upper bounds. Our meaningful choice node is payment method. Perhaps, others should be explored. Buyers and sellers may choose marketplaces on a large number of different margins, such as: market fees, network effects, intensity of competition, encryption methods, etc. These other features of markets certainly bear further research.

Our results add to the current literature on both spontaneous order and reputation systems. Like the Law Merchant, they demonstrate how a marketplace, where feedback mechanisms and reputation are the only things keeping the market functioning, can exist without government regulation. These feedback mechanisms have created an informal institutional framework within which traders exchange goods with confidence (Milgrom *et al.*, 1990). This marketplace demonstrates the shifting institutional structure of black markets in response to new technologies and threats. Silk Road cannabis sales data support the theory that investment in reputation provides a premium to sellers, creating a framework that incentivizes sellers to deliver good service to buyers, despite anonymity and an absence of *ex post* recourse. Reputation's role is especially powerful in this case; it is fundamental to the community's existence.

Because sellers are able to charge premium prices due to their higher relative reputations, this incentivizes them to work to increase their reputation. This particular incentive structure further solidifies the theory that reputation mechanisms are effective. Good reputations allow sellers to make more money and sellers are incentivized to provide quality service to their customers so that they increase their reputation and thus, make higher profits. This supports the Leeson (2005) criteria for an effective signal inasmuch as seller reputation is readily observable, cheap for honest sellers to obtain, and costly for dishonest sellers to garner.

Our results also demonstrate that other factors are certainly at play. Our estimates of β suggest that future revenue streams may not be enough to keep smaller sellers honest. This could reveal itself as a tendency for larger sellers to

dominate the marketplace, or cause other honesty-encouraging mechanisms to emerge. Markets, following the fall of Silk Road, are increasingly outsourcing escrow services. These third-party providers supply different bundles of service; it will be interesting to see which escrow features consumers cluster around.

Further research could analyze this process over a much longer period of time and track certain sellers to see when they enter and exit the market due to reputation. An analysis of the creation of transaction networks, behaviors of entering and exiting sellers, and buyer behavior are also enlightening questions to be explored.

Acknowledgements

The authors wish to thank Peter Leeson, Donald Boudreaux, and Daniel Houser for helpful comments. We also thank Jeffrey Tucker for providing an interview opportunity, and Sam Gelman for his invaluable coding assistance.

References

- Allen, F. (1984), 'Reputation and product quality', *The RAND Journal of Economics*: 311–327.
- Benson, B. L. (1989), 'The Spontaneous Evolution of Commercial Law', *Southern Economic Journal*, 55(3): 644–661. Retrieved from <http://search.proquest.com/docview/56371368?accountid=14541>.
- Briere, M., K. Oosterlinck, and A. Szafarz (2013), 'Virtual Currency, Tangible Return: Portfolio Diversification with Bitcoins', ULB – Universite Libre de Bruxelles, Working Papers CEB: 13–031. Retrieved from <http://search.proquest.com/docview/1438547719?accountid=14541>.
- Chandler, N. (2013), 'How the Deep Web Works', from HowStuffWorks.com website: <http://computer.howstuffworks.com/internet/basics/how-the-deep-web-works1.htm> (accessed July 16, 2014).
- Christin, N. (2012), 'Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace', *Cornell University Library Working Paper*, 12–018.
- Clay, K. (1997), 'Trade without Law: Private–Order Institutions in Mexican California', *Journal of Law, Economics, and Organization*, 13(1): 202–231.
- DarkNet Markets. (2014), Reddit website: <http://www.reddit.com/r/DarkNetMarkets> (accessed October 14, 2014).
- European Central Bank (ECB). (2012), Virtual Currency Schemes, <http://www.ecb.int/pub/pdf/other/virtualcurrencyschemes201210en.pdf> (accessed Mar. 2, 2015).
- Fries, A., R. W. Anthony, A. Cseko Jr, C. C. Gaither, and E. Schulman (2008), *The Price and Purity of Illicit Drugs: 1981–2007* (No. IDA-P-4369). Alexandria VA: Institute for Defense Analyses.
- Gambetta, D. (2010), *Codes of the Underworld: How Criminals Communicate*, Princeton, NJ: Princeton University Press.
- Greif, A. (1989), 'Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders', *The Journal of Economic History*, 49(4): 857–882.

- Greif, A. (1993), 'Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition', *The American Economic Review*, 83(3): 525–48.
- Greif, A. (2012), 'The Maghribi Traders: A Reappraisal?', *Economic History Review*, 65(2): 445–469.
- Grinberg, R. (2011), 'Bitcoin: An Innovative Alternative Digital Currency', *Hastings Science & Technology Law Journal*, 160–207.
- Houser, D. and J. Wooders (2006), 'Reputation in Auctions: Theory, and Evidence from eBay', *Journal of Economics & Management Strategy*, 15(2): 353–369.
- Hu, X., Z. Lin, A. B. Whinston, and H. Zhang (2004), 'Hope or Hype: On the Viability of Escrow Services as Trusted Third Parties in Online Auction Environments', *Information Systems Research*, 15(3): 236–249. <http://dx.doi.org/10.1287/isre.1040.0027>.
- Kilmer, B. and R. L. Pacula (2009), 'Estimating the Size of the Global Drug Market: A Demand-side Approach', *A Report on Global Illicit Drug Markets, 1998–2007*, 25.
- Klein, B. and K. Leffler (1981), 'The Role of Market Forces in Assuring Contractual Performance', *Journal of Political Economy*, 89(4): 615–641.
- Leeson, P. T. (2005), 'Endogenizing Fractionalization', *Journal of Institutional Economics*, 1(1): 75–98. Retrieved from <http://search.proquest.com/docview/213620017?accountid=14541>.
- Leeson, P. T. (2007), 'An-arrgh-chy: The Law and Economics of Pirate Organization', *Journal of Political Economy*, 115(6): 1049–1094.
- Leeson, P. T. (2010), *Anarchy Unbound: How Much Order Can Spontaneous Order Create?*, Cheltenham, UK and Northampton, MA: Elgar.
- McDonald, C. and V. Slawson (2002), 'Reputation In An Internet Auction Market', *Economic Inquiry*, 40(4): 633–650.
- Melnik, M. and J. Alm (2002), 'Does a Seller's E-Commerce Reputation Matter? Evidence from eBay Auctions', *Journal of Industrial Economics*, 50(3): 337–349.
- Milgrom, P. R., D. C. North, and B. R. Weingast (1990), 'The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges and the Champagne Fairs', *Economics and Politics*, 2(1): 1–23.
- Morgan, J. P. (2012), '2012 Online Fraud Report'. https://www.jpmorgan.com/cm/BlobServer/13th_Annual_2012_Online_Fraud_Report.pdf?blobkey=id&blobwhere=1320571432216&blobheader=application/pdf&blobheadername1=Cache-Control&blobheadervalue1=private&blobcol=urldata&blobtable=MungoBlobs (accessed June 25, 2015).
- Mueller, D. C. (1988), 'Anarchy, the Market, and the State', *Southern Economic Journal*, 54(4): 821–830.
- Onion Routing. (2014), Onion Routing website: <http://www.onion-router.net/> (accessed September 7, 2014).
- Posner, R. A. (1980), 'A Theory of Primitive Society, with Special Reference to Law', *The Journal of Law and Economics*, 23(1): 1–53.
- Powell, B. and E. P. Stringham (2009), 'Public Choice and the Economic Analysis of Anarchy: A Survey', *Public Choice*, 140(3/4): 503–538.
- Resnick, P. and R. Zeckhauser, (2001), 'Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System', Mimeo.
- Shapiro, C. (1983), 'Premiums for High Quality Products as Returns to Reputations', *The Quarterly Journal of Economics*, 98(4): 659–680.

Skarbek, E. S. (2008), 'Remittances and Reputations in Hawala Money-Transfer Systems: Self-Enforcing Exchange on an International Scale', *Journal of Private Enterprise*, 24(1): 95–117.

Tucker, J. (2014), Personal Interview, September 23, 2014.

U.S. vs. Ross Ulbricht, Government Exhibit 940D. 14 Cr. 68. New York Southern District Court. 30 Sept. 2013.

Appendix A

Image 1. (Colour online)

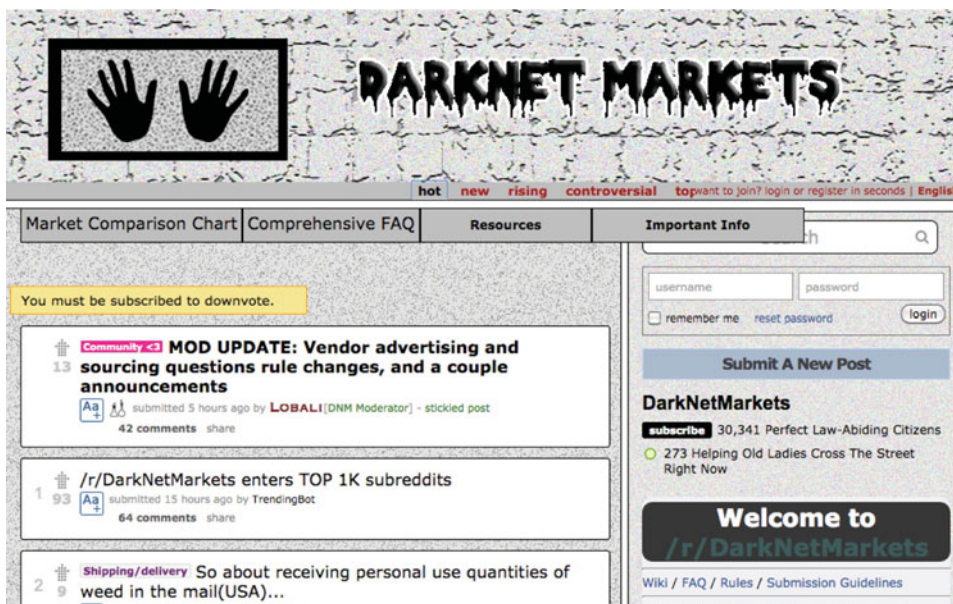


Image 2. (Colour online)




Image 3. (Colour online)

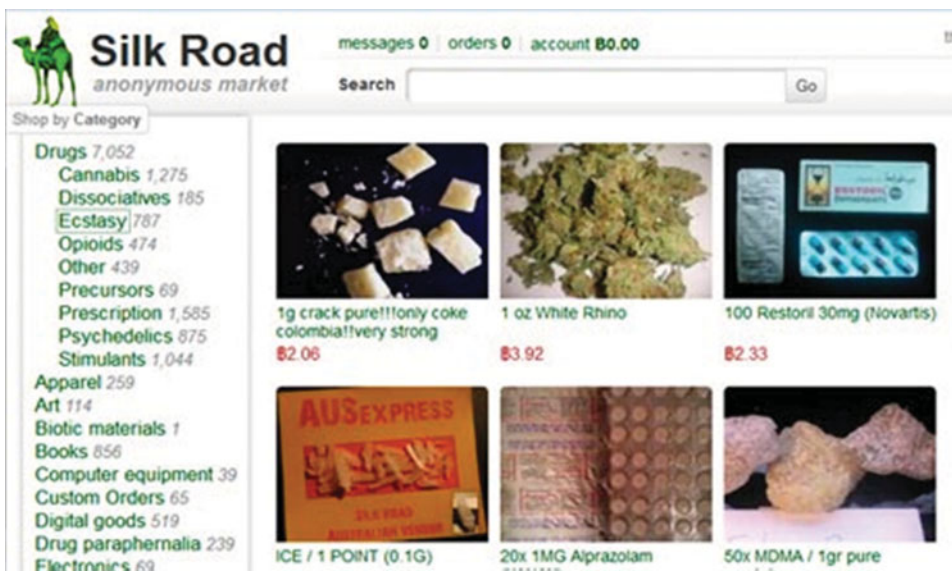


Image 4. (Colour online)

The screenshot shows the Silk Road anonymous market interface. At the top left is the logo with a camel and the text 'Silk Road anonymous market'. To the right, there are navigation links for 'messages 0', 'orders 0', and 'account B0.000'. Below this is a search bar with a 'Go' button. On the left side, there is a vertical menu listing various categories such as Weed, Alcohol, Apparel, Art, Biotic materials, Books, Collectibles, Computer equipment, Custom Orders, Digital goods, Drug paraphernalia, Drugs, Electronics, Erotica, Forgeries, Hardware, Herbs & Supplements, Jewelry, Lab Supplies, Lotteries & games, Medical, Money, Packaging, Services, and Writing, each with a corresponding number. The main content area features a product listing for '10g White Widow'. It includes a photograph of the cannabis plant, a five-star rating with '(688)' reviews, a price of 'B0.262237', and the vendor name 'DutchMagic' with a '93' badge. Shipping information indicates 'ships from: Netherlands' and 'ships to: Worldwide'.

Image 5. (Colour online)

rating	feedback	freshness
★★★★★	Extremely good product for the price, I am a satisfied customer. Will call again.	1 day
★★★★★	awesome	1 day
★★★★★	Perfect	1 day
★★★★★	perfect as usual xxx	1 day
★★★★★	Fuck I was getting really worried! Delivery to Germany 13 days after marked shipped! I was starting to doubt anything in the world. Bit sticky but it will be excellent. Love you DM, sorry for the unnecessary message! efwnlirewtn435ds	1 day
★★★★★	always the same shit!!! THE BEST SHIT YOU CAN GEET HERE ;] 10/10	2 days
★★★★★	good communication with the vendor. not the same hight quality as othertimes but still good	2 days
★★★★★	Quick delivery, excellent stealth	2 days
★★★★★	Received the weed, great grass great stealth, recommended vendor.	2 days
★★★★★	8 Working days to UK - Great Product *****	2 days
★★★★★	5 Days to UK; Fantastic product for price, extrememly kiefy; 1g overweight, quite stalky; Reliable and speedy vendor! 5/5	2 days
★★★★★	My 2nd order, and as before im amazed with this guy :-) 7 days to UK and quality is top standard. I will be back very soon...	2 days
★★★★★	Good Delivery and good smoke	2 days

Image 6. (Colour online)

shipping options

description	est. delivery	shipping price	
Free Shipping (non-tracked mail)	9 days	€0.000000	add to cart

item feedback

30 day average: **4.90**
 60 day average: **4.90**
 Overall average: **4.89**