Reputation Dynamics in a Market for Illicit Drugs

Nick Janetos Jan Tilly

January 2017

WORK IN PROGRESS

Abstract

We analyze reputation dynamics in an online market for illicit drugs using a novel dataset of prices and ratings. The market is a black market, and so contracts cannot be enforced. We study the role that reputation plays in alleviating adverse selection in this market. We document the following stylized facts: (i) There is a positive relationship between the price and the rating of a seller. This effect is increasing in the number of reviews left for a seller. A mature highly-rated seller charges a 20% higher price than a mature low-rated seller. (ii) Sellers with more reviews charge higher prices regardless of rating. (iii) Low-rated sellers are more likely to exit the market and make fewer sales. We show that these stylized facts are explained by a dynamic model of adverse selection, ratings, and exit, in which buyers form rational inferences about the quality of a seller jointly from his rating and number of sales. Sellers who receive low ratings initially charge the same price as highly-rated sellers since early reviews are less informative about quality. Bad sellers exit rather than face lower prices in the future. We provide conditions under which our model admits a unique equilibrium. We estimate the model, and use the result to compute the returns to reputation in the market.

Keywords: online reputation, adverse selection, crime

JEL Classification Numbers: L150, K420

*University of Pennsylvania, 160 McNeil building, 3718 Locust Walk, Philadelphia, PA 19104. Email njanetos@econ.upenn.edu and jtilly@econ.upenn.edu. Thanks to Camilo Garcia-Jimeno, Qing Gong, Ami Ko, Pinar Yildirim, and workshop participants at the University of Pennsylvania for helpful discussion and comments.

1 Introduction

In recent years, online marketplaces for illicit drugs have emerged due to innovations in digital currencies and encryption. In these illegal marketplaces, buyers and sellers transact anonymously and the merchandise is shipped by mail. These marketplaces face significant problems of adverse selection, because based on product descriptions alone, buyers cannot tell good products apart from bad products. To study the role that reputation plays to alleviate the adverse selection problem in these markets, we collect a novel dataset of prices, reviews, and ratings from one such marketplace. In this market, buyers can influence the reputation of sellers by leaving reviews, which affect seller ratings. We document how ratings, prices, and sellers' decisions to exit the marketplace interact.

We then develop a dynamic model of reputation, in which buyers form rational inferences about the quality of a seller jointly from his rating and the number of sales.[1] Sellers who initially receive low ratings charge the same price as sellers who initially receive high ratings, since early reviews are less informative about quality. Bad sellers exit rather than face lower prices in the future. We provide conditions under which our model admits a unique equilibrium. We estimate the model, and use the result to compute the returns to reputation in the market.

Black markets, or markets for illegal goods, are of considerable economic importance, and understudied due to difficulties involved in measuring activity in these markets (see e.g. Levitt and Venkatesh (2000)). Online markets for illicit drugs generate annual revenues in the hundreds of millions of US dollars.[2] Black markets differ from legal markets in a number of important ways. The aspect we focus on in this paper is the role reputation plays in the operation of these markets. A buyer may purchase a good and discover it is of low quality. His only recourse is to then leave a bad review for the seller through a rating system, akin to those used by legal marketplaces such as Amazon and eBay. In legal markets, a buyer may also receive a refund through the market platform itself (both Amazon and eBay have money-back guarantees), or, in extreme circumstances, may take legal action. Neither of these channels are available to buyers in a black market. If lower-ranked sellers receive a lower price, then the rating system works either to enforce cooperation (in the case of moral hazard) or to drive bad sellers out of the market (in the case of adverse selection). In this paper, we focus on adverse selection.

We collect a novel dataset of prices and ratings for illicit drugs from Agora, a popular online marketplace for a vast range of illegal merchandise. We develop a web crawler suitable for the

[1] In the theoretical literature, there are different approaches to modeling reputation. See Mailath and Samuelson (2006) for a survey. We take an adverse selection approach. Sellers differ in their types, which we interpret as qualities. Buyers do not directly observe the quality of sellers. They observe noisy signals, which are informative about the quality of the seller. As in Ekmekci (2011), these signals are modeled as a rating system, in which buyers leave a review after each sale. The reputation of a seller is the market's belief about his type, conditional on the information revealed by the rating system.

[2] The first such online black market was Silk Road, founded in 2011. Soska and Christin (2015) analyze data from Silk Road and impute that the site generated daily revenues of about $300,000 in early 2013, implying annual revenues in excess of $100 million. Later in 2013, the Silk Road's founder was arrested and the dark net market shut down. A variety of other market platforms then imitated Silk Road, and some of these market platforms continue to operate. They comprise a relatively small part of the global drug trade, where global revenue is estimated to be in the hundreds of billions of US dollars.

dark net and use it to archive listing pages, vendor pages, and reviews. The dataset covers the time period from January 2014 until the market's closure in August 2015. Using this dataset, we document the effect a seller's rating has on the prices he charges.

In line with the prior literature that studies legal markets, there is a positive relationship between the price and the rating of a seller (for surveys, see Dellarocas (2002) and Cabral (2012)). We show that this effect is increasing in the number of reviews left for a seller. A mature highly-rated seller charges a 20% higher price than a mature low-rated seller. For young sellers, the rating has no relationship to the price. Sellers with more reviews charge a higher price than sellers with a low number of reviews regardless of rating. Our findings are consistent with adverse selection and buyers who draw rational inferences from ratings: When a seller only has a few reviews, buyers do not draw much inference about the quality of a seller from the rating, but when a seller has many reviews, the rating is more informative, and has a stronger effect on the price. Bad sellers exit the marketplace rather than face lower prices in the future.

We introduce a dynamic model of adverse selection and ratings that is consistent with the data. In our model, sellers enter the market, and are assigned a type, unobserved by buyers. As sales are made, buyers sample the goods, and can leave feedback with the market. We are agnostic about the exact process through which buyers form opinions and leave feedback. Following prior theoretical literature on rating systems (see e.g. Ekmekci (2011); Bohren (2012); Hu (2015)), we assume only that the market offers some type-dependent, stochastically evolving, publicly available information about each buyer. In our case this information consists of the rating and the total number of sales made by a seller.

Buyers are short-lived and do not observe the full history of the seller's public characteristics, only the public rating at the time at which the buyer enters the game. From this information, buyers form inferences about the quality of sellers consistent with the sellers' equilibrium exit strategies, and sellers are paid a price for their product which depends on the inference buyers draw in equilibrium about seller quality. Our model differs from the existing theoretical literature on rating systems by considering a problem of adverse selection; the sole strategic decision made by a seller is whether or not to exit the game. For some parameter values there may be multiple equilibria. We establish restrictions on the discount rate so that the model admits a unique equilibrium.

In our model, the market's rating system acts to separate low and high quality sellers, alleviating adverse selection in the market. Low quality sellers who are identified as such prefer to exit the market. This induces a tension between exit and the rating system—if, for example, it were the case that low quality sellers strictly preferred to exit at a low rating, then rational buyers should infer that low rated sellers are in fact high quality, meaning that *low* rated sellers charge a *high* price in equilibrium. At the same time, exit of low-rated sellers introduces survivorship bias into reduced form estimates of the relationship between price and rating. These two channels—inference on the part of consumers, and survivorship bias—both bias reduced form estimates of the relationship between price and rating downward. We show by example that these two effects may be so large that it may happen that some sellers receive a *higher* price if rated poorly, even as the returns to reputation are positive.

Prior literature

The work in this paper contrasts with previous research in two ways: First, we consider a novel market platform, in which contracts cannot be enforced. Second, we do not solely focus on the impact a seller's rating has on his price, rather, we focus on the joint role of ratings, exit, prices, and inferences buyers draw about the quality of sellers.

Perhaps the closest paper to ours is Saeedi (2014), who estimates a structural model of adverse selection and ratings on eBay. Our approach differs in the following key ways: First, we consider a richer set of publicly observable seller characteristics, including the rating, a feature made available by the relatively more disperse ratings in our dataset compared to eBay. Second, we model price formation differently. Saeedi assumes that buyers draw no inferences from the price. In contrast, we assume that sellers are price takers and paid their expected worth conditional only on publicly available characteristics.

Prior work on the market for illicit drugs has focused mainly on the determinants of prices, using the STRIDE dataset which records prices of seized illicit drugs in the United States. Using these data, Galenianos et al. (2012) and Galenianos and Gavazza (2015) focus on learning under moral hazard, but do not focus on rating systems. Dobkin and Nicosia (2009) look at the effect of government seizures of methamphetamine precursors on the price of methamphetamines.

Our findings are consistent with the literature on *legal* online markets. A small but growing literature attempts to estimate the 'returns to reputation', i.e. the additional value that a highly-rated seller receives in a market over a low-rated seller (see Ederington and Dewally (2002), Melnik and Alm (2002), Resnick and Zeckhauser (2002), Eaton (2005a), Eaton (2005b), Houser and Wooders (2006), Cabral and Hortaçsu (2010), Yoganarasimhan (2013), Saeedi (2014), and Jolivet et al. (2016)). In almost all of the existing literature, a statistically significant and positive, but small and fragile effect of reputation on price is found.

The rest of this paper is structured as follows. In Section 2, we discuss the institutional details of these markets and provide some historical background. In Section 3, we explain how the data were collected, and establish some stylized facts about the role reputation plays in this market. In Section 4, we present a dynamic model of reputation consistent with these stylized facts. In Section 5, we estimate the model. In Section 6, we use the estimated model to compute the return to reputation in this market. We conclude in Section 7.

2 Background: Dark Net Market Places

The first online black market was Silk Road, founded in 2011 and shut down by law enforcement 2013. Among its successors is the Agora marketplace, which is the data source for this paper. The Agora marketplace launched in late 2013, and went on to weather two of the major shocks to dark net markets. First, in November 2014, the FBI seized the servers of many major markets as part of ‘Operation Onymous’. The Agora marketplace was not affected, and subsequently became one of the major platforms for trading on the dark net. In March 2015, one of Agora’s competitors, the Evolution marketplace, closed suddenly, taking the assets held in escrow along with it, after which the Agora marketplace became the largest market platform in the dark net. Traffic spiked, and several months later, in August 2015, the administrators of the site announced that due to security concerns, they were temporarily taking their operations offline. Funds were returned to sellers and buyers, and the market servers were shut down.

While it was operational, the Agora marketplace was (like most dark net marketplaces, see Figure 1) run in a way recognizable to any user of Amazon. Sellers could, after paying a fee (around $500), open an account and create listings. Sellers were free to choose the price for each listing. Buyers could then see a list of goods for sale using a search tool or by clicking on a category. Each entry for a listing contained its price, the number of sales made by a seller, the net rating of the seller, and information on the shipping source and destination of the goods. Potential buyers could then click to see more details about a listing, such as a more detailed description written by the seller, as well as individual reviews left by buyers.

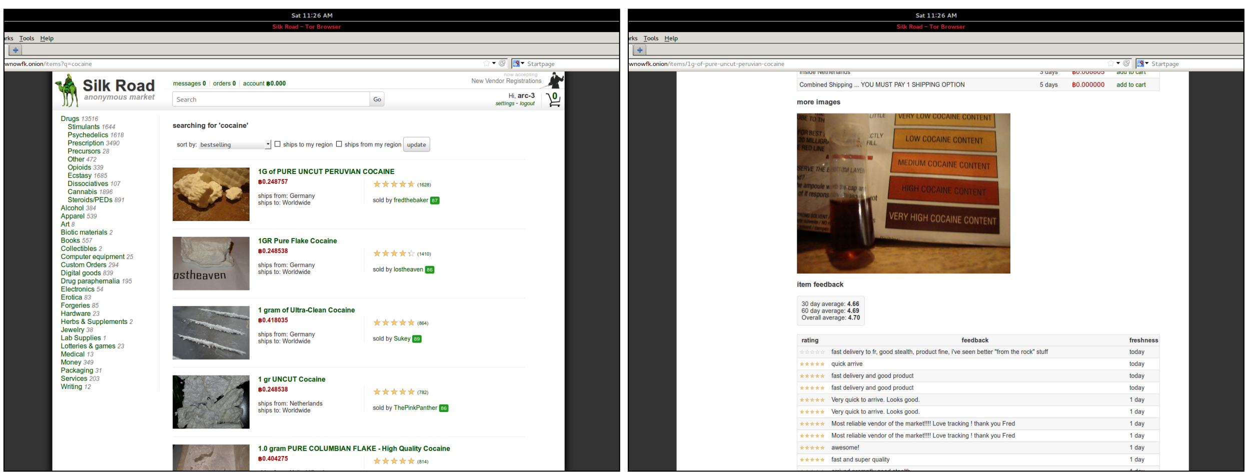

Figure 1: Screen shots from Silk Road, mid-2015.

Note: On the left, we show a screen shot from searching for ‘cocaine’ on Silk Road. On the right, we show a screen shot of the reviews page of a seller, as well as a photo provided by the seller illustrating the high cocaine content of his product.

Once a buyer had made a decision, they could click to add the listing to a shopping cart. All items were denominated and purchased using bitcoins. To pay, a buyer would transfer their bitcoins directly to the marketplace itself. Many of the marketplaces offered a bitcoin ‘tumbler’, a bitcoin-laundering algorithm which transfered the bitcoin between many different addresses to anonymize its source location.[3] The Agora marketplace did not offer a bitcoin tumbler, rather, participants were encouraged to anonymize their own currency.

Once a transaction had been requested, the marketplace would check to make sure that a buyer

[3]Bitcoins are a non-fungible currency, and in theory it is possible to track its history through every transaction made.

had the requisite number of funds held in escrow by the marketplace. A postal address was sent to the seller, who would then ship the item to the buyer using regular mail. Items were usually disguised to not draw suspicion from the postal service or law enforcement. Once the item arrived, buyers were responsible for returning to the marketplace and marking the transaction as having completed, at which point, the marketplace would release the funds held in escrow to the seller. If the buyer did not mark the transaction as complete, the funds would be released after several weeks, unless the buyer decided to contest the sale, in which case the Agora marketplace offered mediation services between buyer and seller. We do not observe instances of this occurring, and to our knowledge this was a rare outcome. Some sellers preferred to bypass the escrow system, advertising themselves as sellers who required buyers to finalize the transactions early, releasing the funds before the goods had actually been received.

3 Data

3.1 Data Collection and Summary Statistics

We scrape data from the Agora marketplace at weekly intervals from January 2014 until the market's closure in August 2015. We develop a web crawler to archive vendor pages, listing pages, and reviews.[4] Vendor pages contain the market-reported aggregate rating and the number of sales made. These vendor pages change frequently as vendors make additional sales. Listing pages contain detailed production information, including a title and description of the product, its price, a picture, and whether the vendor requires buyers to pay up-front or after arrival of the shipment. Note that the listing page may change over time, because vendors can edit all the information on the listing page. Listing pages also contain the entire history of reviews. For each review we observe and record

1. the date at which a review was left,
2. the price of the good at the time the review was left,
3. the rating which was left for the review, which is a number between 0 and 5,
4. the market-reported aggregate rating for the vendor for whom the review was left,
5. the sort of good on sale, in this paper we focus on cannabis, MDMA, heroin, and cocaine, and
6. various vendor and product characteristics, such as the total sales made by that vendor and whether he required buyers to pay up-front or after the shipment arrived.

[4] A web crawler is a computer program which visits a link, records all the links it finds there, then visits each of those links, recursively repeating the process in order to visit every page of a website. Previous research archiving the dark web uses infrequent archives of the entire website to develop a 'snapshot' of the site at different times. Since sites on the dark web are often inaccessible due to network problems, we develop a bespoke web crawler which continually visits the site, cross-referencing the information it finds with previous visits to produce a high-frequency history of the site.

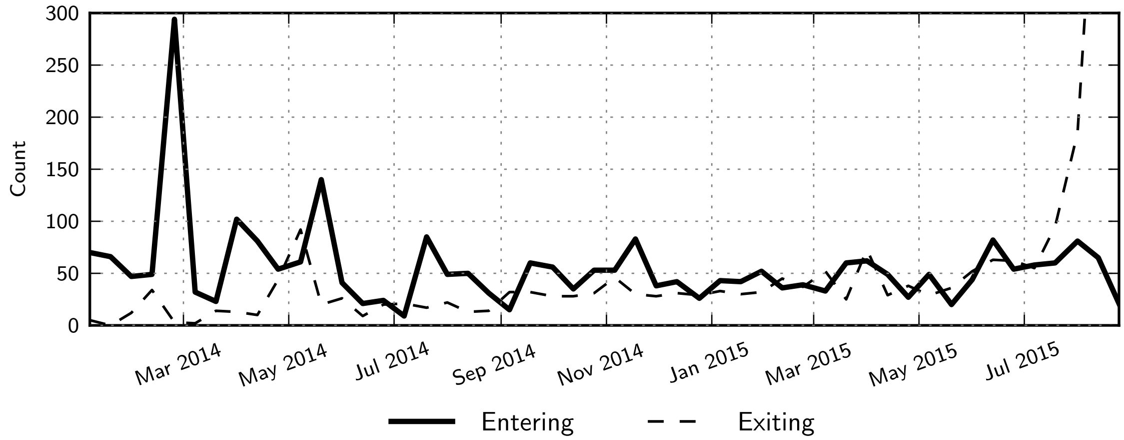

Figure 2: Distribution of entry and exit by vendors over the course of the market lifespan.

Note: The figure shows the number of vendors who enter the market (solid line) and exit the market (dashed line) over time. Entry is defined as the date the first review was observed in the sample, and exit as one week before the last review observed in the sample. The numbers of entries and exits are relatively stable throughout the lifetime of the market. Due to our definition of exit, the number of exits spikes when the market closed.

For the four product categories that we focus on, cannabis, MDMA, heroin, and cocaine, we collected information on $1,482$ vendors, $14,292$ listings, and $572,266$ reviews.

We do not observe transactions directly, but we find that the total number of reviews is a good approximation for the total number of sales made by each particular vendor. Therefore, we treat each review as a transaction and we will occasionally use the terms review and transaction interchangeably. We do not observe the time of sale, only the time at which the review was left. Therefore, the price at the time of sale is imputed by using the price one week before the review was left.

The entry and exit date of vendors is not directly observed. Entry is not directly observed, because some vendors entered the market before we begin scraping the platform. Exit is not directly observed, because exit is indistinguishable from inactivity. We therefore use reviews—our proxy for transactions—to determine entry and exit. A seller is defined to have entered one week before the first review is left, and exited one week before the last review observed in the sample. See Figure 2 for the distribution of entry and exit dates through the lifespan of the market. Entry and exit rates are relatively homogeneous throughout the lifetime of the market, with the exception that many sellers exited when the market closed.

The Agora market does not have an explicit mechanism for registering the sort of good on sale. Goods are categorized (e.g., `MDMA---Pills`), but the sort of item on sale and the quantity on sale is not explicitly registered with the site. To address this, we focus on homogeneous goods and extract quantity information from the listing text by exploiting conventions developed on the marketplace for listing quantities. For example, a typical listing might be titled

```
XTC pills---best quality---fresh from Amsterdam---120mg x5,
```

from which the scraper extracts the relevant information that this is a listing for ecstasy ('XTC') pills, each containing 120 milligrams of MDMA, sold in batches of five. The prices are then normalized to be in dollars per gram, and consistency checks are run on price data to ensure that prices found are broadly consistent with other data.

We choose to focus on the top four categories of goods sold: Cannabis, MDMA, heroin, and cocaine. We report summary statistics for these four categories in Table 1. In total we observe 1,482 vendors who sell in any of these four product categories. Among them, 1,080 vendors exclusively sell in one product category, 95 in two, 298 in three, and only four vendors sell in all four product categories. On average, each seller has received 386 reviews across the four product categories. The average lifespan of a vendor ranges from approximately 300 to 350 days, where the lifespan per product category refers to the number of days between the first and last posting of a product in a particular category. A little over one third of all vendors was still active when the market closed in August 2015. This is roughly equal across product categories. Cannabis is by far the cheapest product per gram (approximately 12 dollars per gram), Heroin the most expensive (approximately 162 dollars per gram). For all products, the ratings that vendors receive are very high. The average rating across product categories is 4.93 out of 5.00. Per product category, the average ratings range from 4.80 for Heroin to 4.97 for MDMA.

	All	Cannabis	MDMA	Heroin	Cocaine
Number of vendors	1,482	824	508	192	473
Number of reviews per vendor	386.1	338.7	236.2	288.7	249
Average life span per vendor (days)	349.1	344.3	337.4	307.8	319.2
Share of vendors active when market closed	0.38	0.35	0.33	0.43	0.33
Average price per gram	49.83	12.11	37.45	162.32	98.86
Average rating	4.93	4.94	4.97	4.80	4.92

Table 1: Summary Statistics

Note: The table reports summary statistics for the four product categories that we consider: Cannabis, MDMA, heroin, and cocaine. The number of vendors across product categories does not sum to the total number of vendors, because some vendors are active in multiple product categories. Share of vendors active when market closed refers to the share of vendors that we observe making sales after August 1, 2015 (the market was shut down in late August). Prices are reported in USD per gram, where we obtain the dollar amounts by converting bitcoin to U.S. dollars at the time of the transaction.

Figure 3 plots the volume of reviews left weekly for each of these four categories. The figure shows that the market's traffic steadily increased over its lifespan, reaching its peak after the Agora marketplace's main competitors were shut down during 'Operation Onymous'. In the last months of its life, the Agora marketplace began to experience frequent periods during which it was inaccessible due to security precautions taken by its owners. The marketplace was eventually was shut down, also due to security concerns.

Figure 4 shows the derived price at the imputed transaction date, in dollars per gram, for each of the four product categories. Since prices are denominated in bitcoins, the value of which fluctuated wildly for the duration of the sample, we use the market-reported exchange rate between bitcoins

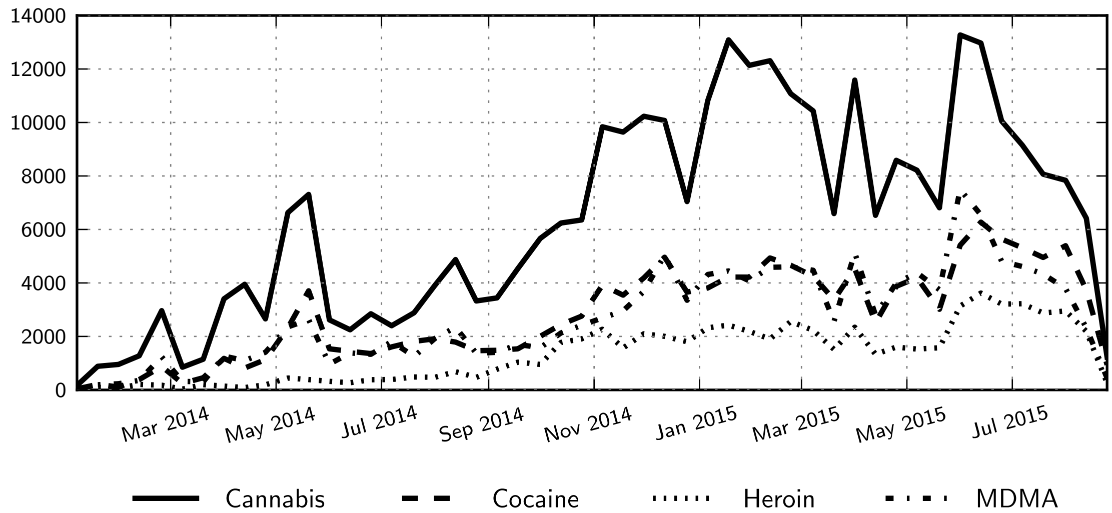

Figure 3: Volume of reviews left, over time, by category.

Note: The figures show the total number of reviews per week that were left for products that fall into the categories cannabis, MDMA, heroin, and cocaine over the entire lifespan of the marketplace. We use reviews as a proxy for transactions. Figure 3 shows that the market's traffic steadily increased over its lifespan, reaching its peak after the Agora marketplace's main competitors were shut down during 'Operation Onymous'.

and US dollars to convert all prices to US dollars. The price distributions or all product categories are humpshaped with a fair degree of dispersion. Overall, the price distributions look reasonable giving credence to our classification algorithm to obtain products, quantities, and weights as well as the currency conversion.

As a point of reference, Figure 5 shows the distribution of prices for Heroin and Cocaine in the STRIDE dataset (DEA, 2016). The STRIDE dataset is provided by the United States Drug Enforcement Administration with information on prices of drugs that were seized between 2007 and 2013. We obtained the full STRIDE dataset through a Freedom of Information Act Request. The price distributions of Heroin and Cocaine—the other product categories, Cannabis and MDMA, are not part of the STRIDE dataset—look similar to the price distributions that we report from the online marketplace. In particular, they are similarly disperse. On average, prices in the online marketplace are higher. We can only speculate why prices in the online marketplace are higher. It is perceivable that vendors in online markets charge a convenience premium. Also, the STRIDE data is a selected sample, since it only contains data from seizures, not from transactions.

Vendors differ in experience and success in the market. See Figure 6 for distributions of sales volume and final age of vendors by the end of sample. Many vendors entered, made a few sales, and then left the market. Other vendors stayed for the full duration of the market of approximately 600 days.

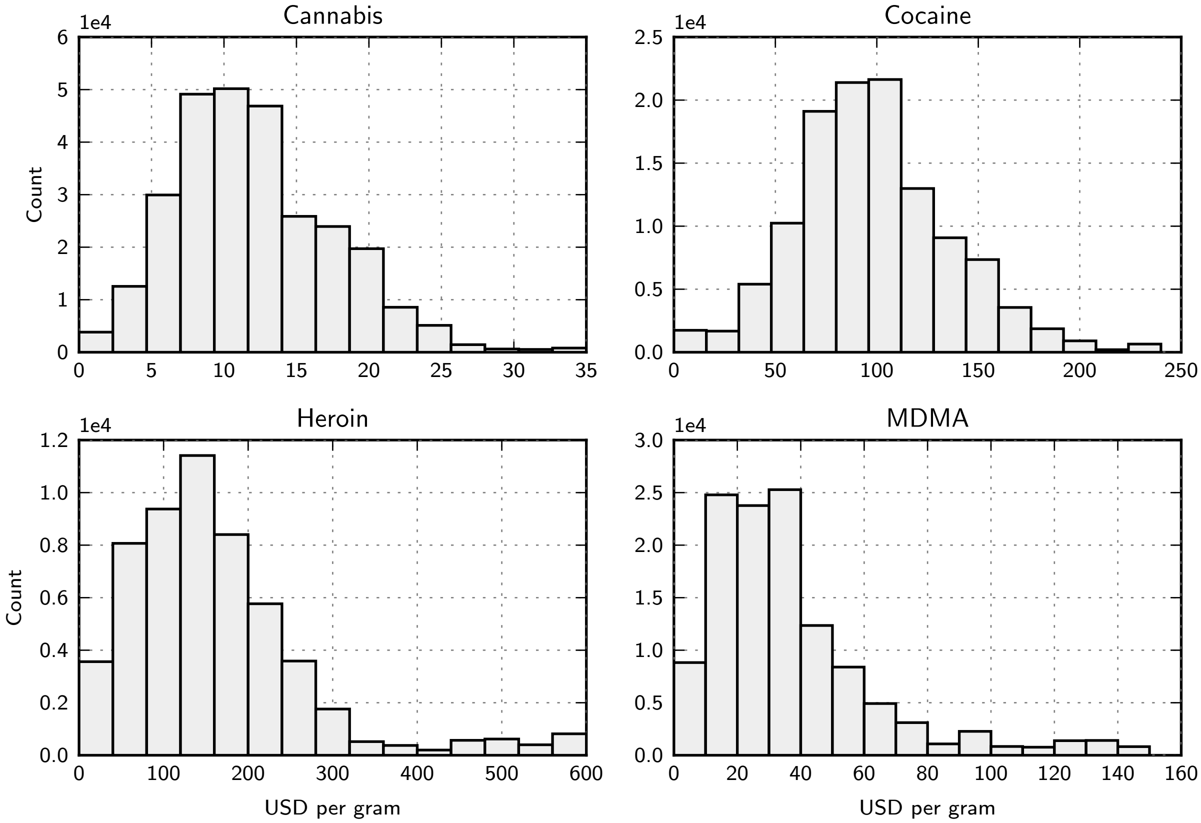

Figure 4: Distribution of the prices of illicit drugs in the Agora market.

Note: The figures show the distribution of prices for cannabis, MDMA, heroin, and cocaine over the entire lifespan of the marketplace. The unit of observation is a review, which we use as proxy for transactions. The y-axes are denoted in $10,000$. The x-axes are denoted in USD per gram, where we obtain the dollar amounts by converting bitcoin to U.S. dollars at the time of the transaction. The price distributions of cannabis, MDMA, heroin, and cocaine are hump-shaped. The price distribution for Heroin has a long right tail.

3.2 Stylized Facts

In this section, we establish the following stylized facts from the data about the relationship between a seller's price, rating, and total reviews left.

1. There is a positive, significant, but small relationship between the price a seller charges, and a seller's rating. Moving from the 25th percentile to the 75th percentile in ratings is associated with a 0.05% increase in price.

2. The more reviews left for a seller, the greater the relationship between a seller's price and his rating. For the largest sellers in the sample, moving from the 25th percentile to the 75th percentile in ratings is associated with a 19.7% increase in price.

3. There is a positive relationship between number of reviews left for a seller and the seller's price. Sellers just entering the market charge, on average, 10% less than mature sellers with more than 5,000 reviews.

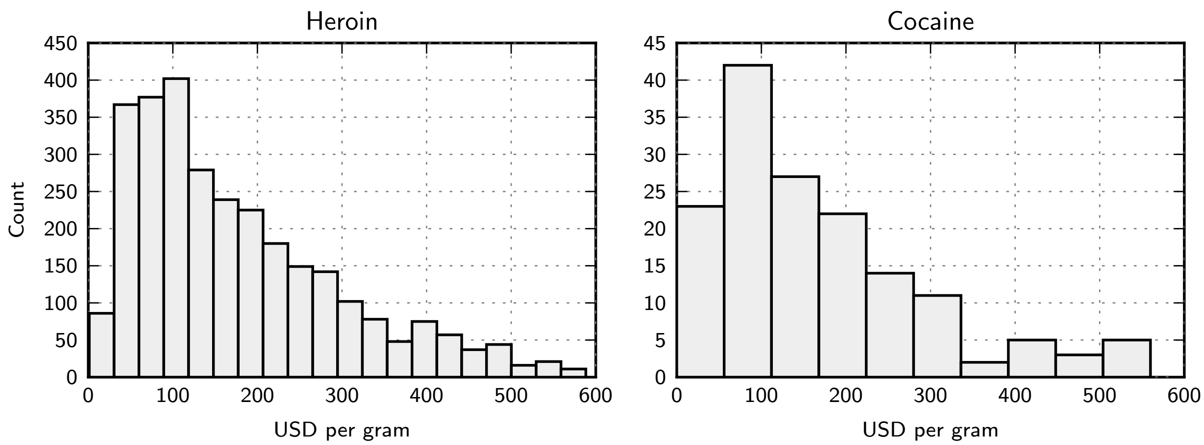

Figure 5: Distribution of the price of heroin and cocaine in the STRIDE dataset

Note: The figures show the distribution of the price of heroin and cocaine in the STRIDE dataset. The STRIDE dataset is a dataset with information on prices of drugs that were seized between 2007 and 2013. We obtained the full. The distribution of prices in the STRIDE dataset displays a spread and mean consistent with the prices we find online.

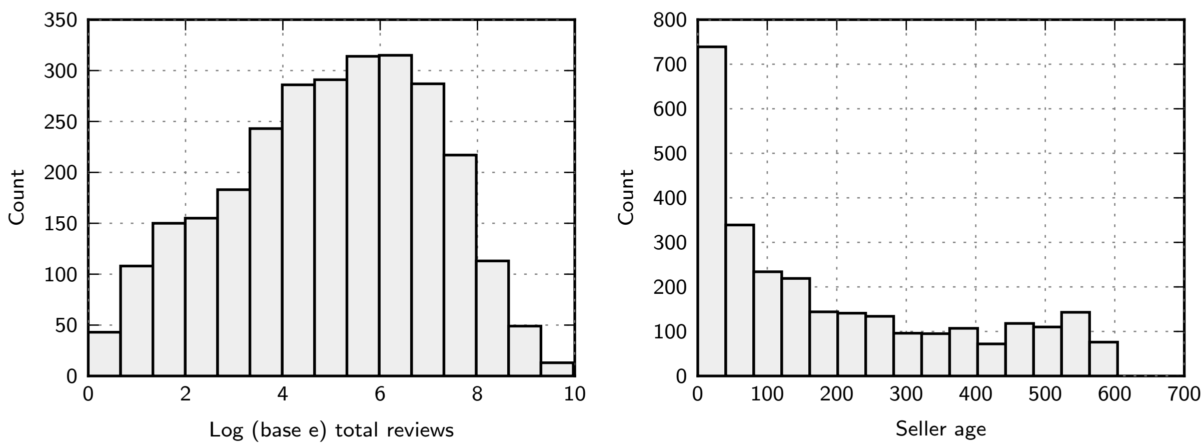

Figure 6: Distribution of seller size and age.

Note: The figure on the left shows the distribution of the log of total sales per seller (proxied by the number of reviews left), where we again restrict our analysis to products in one of the four main categories (cannabis, MDMA, heroin, and cocaine). The figure on the right shows the distribution of the final age (in days) of vendors in the sample for vendors who have sold products in one of the four main categories that we focus on.

4. The more reviews a seller receives, the smaller the impact of an additional review on his price.

In addition, we look at patterns of exit as a function of rating and total reviews left.

5. A seller's lifespan can be broken into three stages, an entry stage, an exit stage, and a mature stage.

6. The lower a seller's rating, the more likely he is to exit.

	(1)	(2)	(3)	(4)
Dependent variable	log(p)	log(p)	log(p)	log(p)
Rating	0.10***		-0.03***	
	(0.00)		(0.00)	
Rating (large sellers)		1.80***		-0.25***
		(0.03)		(0.01)
Rating (small sellers)		0.07***		-0.01***
		(0.00)		(0.00)
Total reviews (1,000s)	0.03***	0.03***	0.00***	0.00***
	(0.00)	(0.00)	(0.00)	(0.00)
Quantity (g)	-0.15***	-0.15***	-0.11***	-0.11***
	(0.00)	(0.00)	(0.00)	(0.00)
Year	-0.02***	-0.03***	-0.02***	-0.01***
	(0.00)	(0.00)	(0.00)	(0.00)
Age (years)	0.02***	0.01*	0.00***	0.00***
	(0.00)	(0.00)	(0.00)	(0.00)
Total sales (1,000s)	-0.05***	-0.04***	0.01***	0.01***
	(0.00)	(0.00)	(0.00)	(0.00)
Intercept	0.03***	0.03***	0.01***	0.01***
	(0.00)	(0.00)	(0.00)	(0.00)
R^2	0.09	0.10	0.08	0.08
N	280366	280366	280366	280366
Fixed effects	No	No	Yes	Yes

Table 2: Regressions of log price residuals on rating, quantity, and sales.

Note: The table shows the coefficient estimates from regressions of the log price on rating, quantity, and the number of sales made. The data consist of vendors who sell cannabis, but similar results hold for the other categories of products. The unit of observation is a review left, our proxy for the number of transactions. Column (1) reports the estimates of coefficients in equation (1) for the entire sample. Column (2) estimates the coefficient on ratings separately for large (number of transactions >50th percentile) and small sellers. Columns (3) and (4) repeat the exercise, but with seller fixed effects.

Price, rating, and total reviews

We begin by analyzing the relationship between prices, ratings, and total number of reviews. We replicate the work of prior literature that studies legal markets and look for reduced form evidence of the impact of the rating on prices. We estimate the model

$$\log(\text{Price}_{it}) = \beta \times \text{Rating}_{it} + \gamma \times \text{Other product characteristics}_{it} + \eta_i + \varepsilon_{it}, \quad (1)$$

where i refers to a particular seller, and t to a particular transaction. The coefficient β captures the effect of rating on price. γ is the coefficient vector associated with additional controls, which include factors such as dummies for the locations vendors ship from and to, the type of product on sale, and age of the vendor and time trends. η_i are vendor fixed effects and ε_{it} is a disturbance term.

The first column in Table 2 displays the estimates of the coefficients β and γ. The data consist of vendors who sell cannabis, but similar results hold for the other categories of products. Consistent with previous literature, a positive, statistically significant, but small effect is found of the rating on the price. We do not interpret the coefficient estimates in Table 2 as causal. Rather we interpret them as conditional means. To interpret the point estimate of 0.10, a seller at the 25th percentile in ratings (around 4.90/5) charges an average 1% lower price than a seller in the 75th percentile in ratings (or 5/5).

The small size of the relationship between rating and price is puzzling, because a higher rating, all else equal, should be a signal of a more valuable product to consumers. One resolution to the puzzle is unobserved seller heterogeneity (see e.g. Resnick and Zeckhauser (2002)). However, we see all relevant seller characteristics. Since sellers are anonymous in our data, their seller page contains all information available to buyers. This is contrast to legal marketplaces, where additional information about sellers might be available elsewhere.[5]

We favor a different explanation for the small size of the relationship between rating and price. Intuitively, the rating of a seller is only informative about the true quality of the seller, once the seller has accumulated a large number of reviews. A seller with a high rating and only a small number of reviews may have just got lucky. A seller with a high rating and a large number of reviews must in fact be delivering a high quality product.

To test this, the second column in Table 2 reports the estimate of β when we allow the coefficient on ratings to depend on whether a seller falls in the top or bottom 50% of total sales made. For sellers in the top 50% of sales made, the relationship between price and rating strengthens to 1.80. To interpret this, a seller moving from the 25th to the 75th percentile in ratings charges an average 19.2% more. This fact suggests that the impact of a bad rating on the price increases as more reviews are left.

We further highlight this point in Figure 7. We report the coefficient estimate of ratings on price, β, broken down by different quantiles of total reviews left. The effect of rating on price is zero when only a small number of reviews have been left, but becomes much larger as the number of reviews increases. For example, after 10,000 reviews, moving from the 25th to the 75th percentile in ratings in the market for cannabis is associated with a 35% higher price, a large and statistically significant number.

The above supports the hypothesis that buyers treat reviews of sellers as informative signals of the quality of the seller, and draw inferences about seller quality accordingly. When a seller does not have many reviews, the rating is uninformative about his quality and has no effect on price. When a seller has many reviews, the rating is very informative and has a large effect on price. Not only does the effect of rating on price increase as sellers receive more reviews, but a sellers with more reviews also charge higher prices. For example, a seller with 10,000 total reviews charges on average 35% more than a seller with only one review.

[5] Newberry and Zhou (2016) investigate this possible explanation for Alibaba's Tmall, China's largest business-to-consumer marketplace. They find that the impact of ratings decreases with the offline presence of a seller.

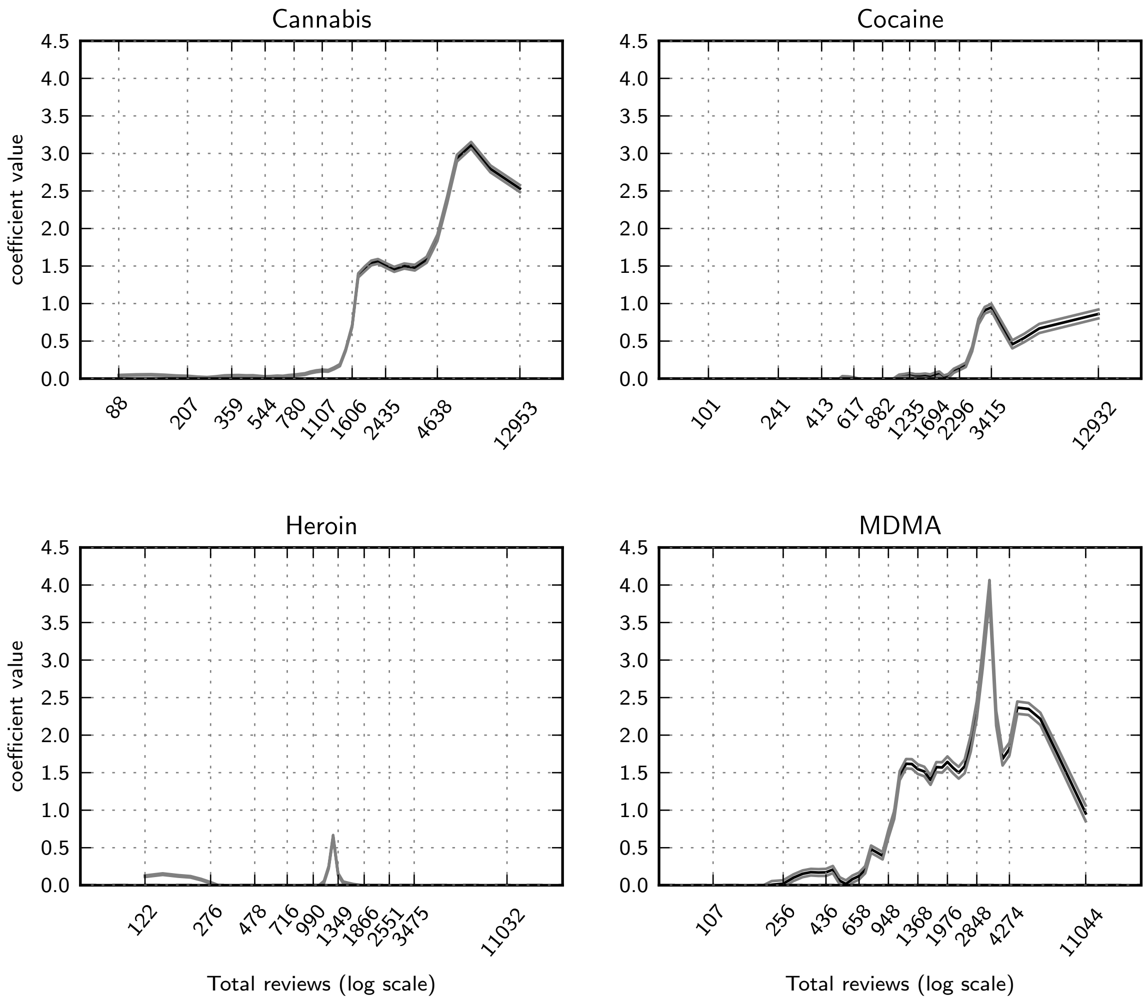

Figure 7: Effect of rating on log price broken down by number of reviews.

Note: We show the regression coefficient β (and their 95% confidence intervals in gray) from the linear regression model shown in equation (1) broken down by different quantiles of total reviews left. Total reviews are denoted with a log scale. The effect of rating on price is zero for sellers with few sales. The effect is large and positive for sellers with many sales. The product category 'Heroine' is an outlier.

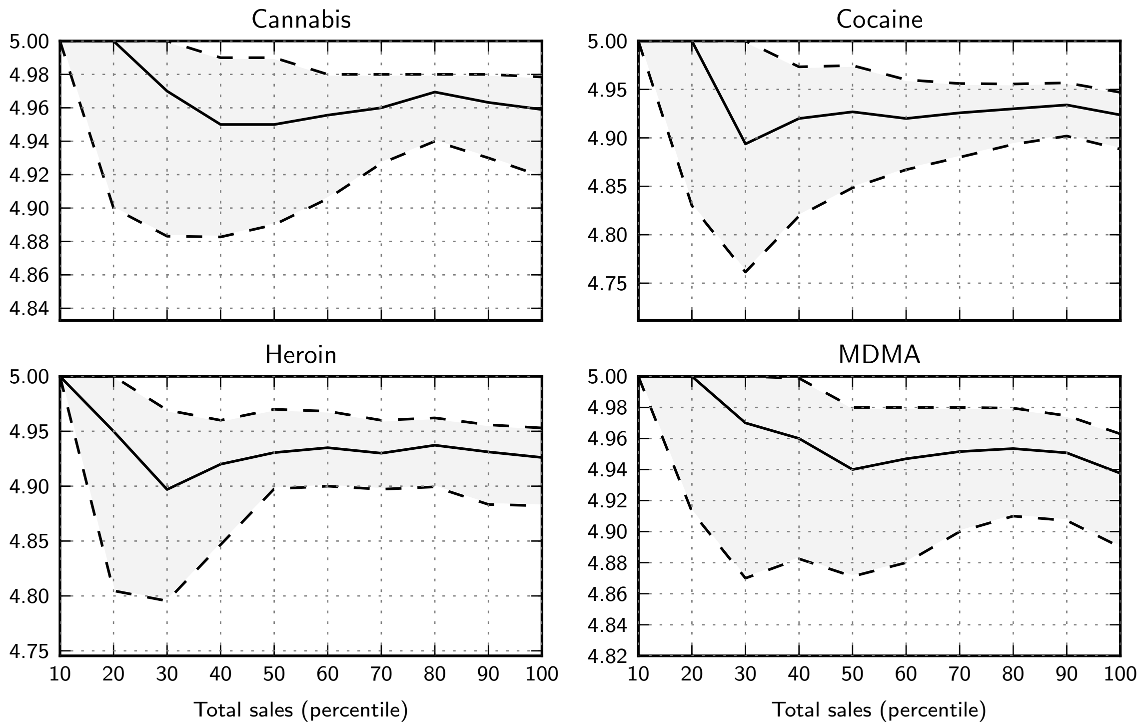

Figure 8: Dispersion of ratings, as a function of sales made.

Note: The figures show a seller's rating as a function of the total number of sales made. Solid line represents the median rating, dashed lines indicate the 30th and 70th percentile of ratings and give an indication of the variance of ratings. In the 0th–10th percentile for sales made (the entry stage), most sellers have the top rating of 5. In the 10th-30th percentile for sales made, sellers ratings spread apart as more reviews arrive. In the 30th-50th percentile (the exit stage), low-rated sellers exit, rather than sell at lower prices, resulting in a narrowing of the rating spread. Finally, in the 50th-100th percentile (the mature stage), the spread and median rating stabilize.

The results of this section suggest that the absence of a strong effect of rating on price is due to a composition effect. For young sellers with few reviews, the rating has little bearing on the price. For older sellers with many reviews, the rating has a substantial effect on prices.

Exit, rating dispersion, and total reviews

The composition effect is amplified by the fact that poorly rated sellers are more likely to leave the market. We investigate the co-evolution of ratings dispersion, ratings, total reviews, and exit.

In Figure 8, we plot the distribution of seller's average weekly rating. Due to granularity in the ratings left (most sellers are in the 4.90–5.00 range, meaning a total of 10 possible ratings, because average ratings are reported with two decimal points), we smooth ratings by using weekly averages. The average rating of a seller tends to follow a 'U'-shape, initially dipping before leveling up again. As it dips, the variance of ratings in the population increases, and as it rises back up, the variance decreases. What causes these U-shaped dynamics? We argue that the answer is two-fold: The design of the rating system itself accounts for the initial dip, while survivorship bias accounts for the subsequent rise.

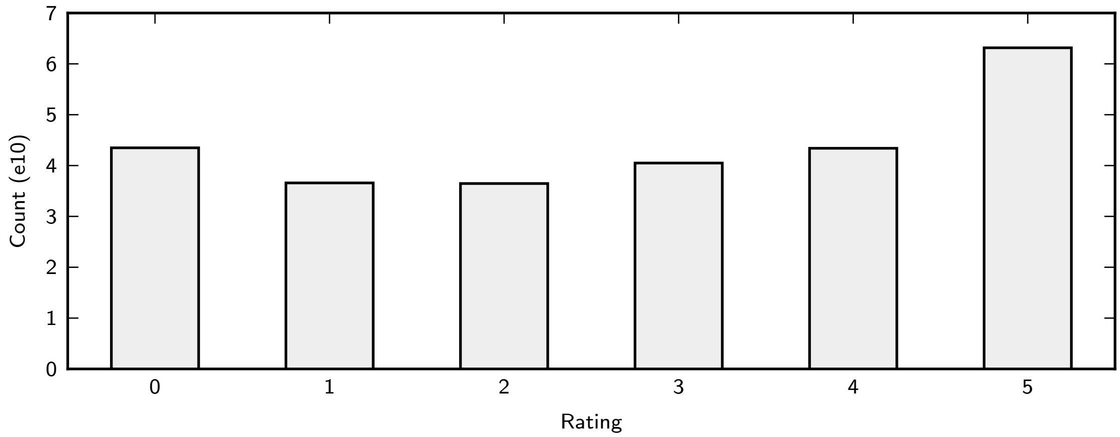

Figure 9: The distribution of ratings received by sellers.

Note: The distribution displays the 'J' shape characteristic of online ratings: Most ratings are 5/5, 4/5, or 0/5. Note the log scale—the number of 5/5 ratings is two orders of magnitude above the number of 0/5 ratings

To see why the design of the rating system causes the initial dip, note that most reviews result in the maximum rating possible, 5/5, the next-most number of reviews result in the worst rating possible, 0/5. As a result, only a small fraction of sellers observed to be entering the market have anything less than the maximum possible rating of 5. Therefore, the initial drop in average rating, and corresponding rise in rating dispersion, as seen in Figure 8 in percentiles 10–30, is a consequence of the fact that the vast majority of sellers enter with a perfect rating of 5/5, and so all sellers see their rating drop as they begin to make more sales.

After the rating has fallen and the rating dispersion increases, the average rating rises and the rating dispersion falls. We argue this is a consequence of survivorship bias. The sellers most likely to continue selling in the market are the high-rated sellers. Hence, past some point, the average rating is increasing and the rating dispersion decreasing, as low-rated sellers exit.

Interpretation

In the remainder of the paper, we argue that these stylized facts are consistent with a dynamic model of ratings, exit, and adverse selection. In the model, when sellers enter the market, they draw a permanent quality type. Based on this type, they then decide in subsequent periods whether or not to exit the market. We show how in some equilibria, the seller's life-cycle may be broken into four stages. See Figure 11 for a graphical illustration, and compare to Figure 8 for the corresponding patterns in the data.

Initially, entering sellers receive the highest possible rating for most reviews. In the entry stage, their rating is an almost entirely uninformative signal of quality. As more reviews accumulate, buyers move into the young stage. Here, the dispersion of ratings increases rapidly, and the average rating falls for all sellers, resulting from the fact that all sellers began at the highest possible rating. Faced with the prospect of lower prices at low ratings, sellers move into the exit stage, in which

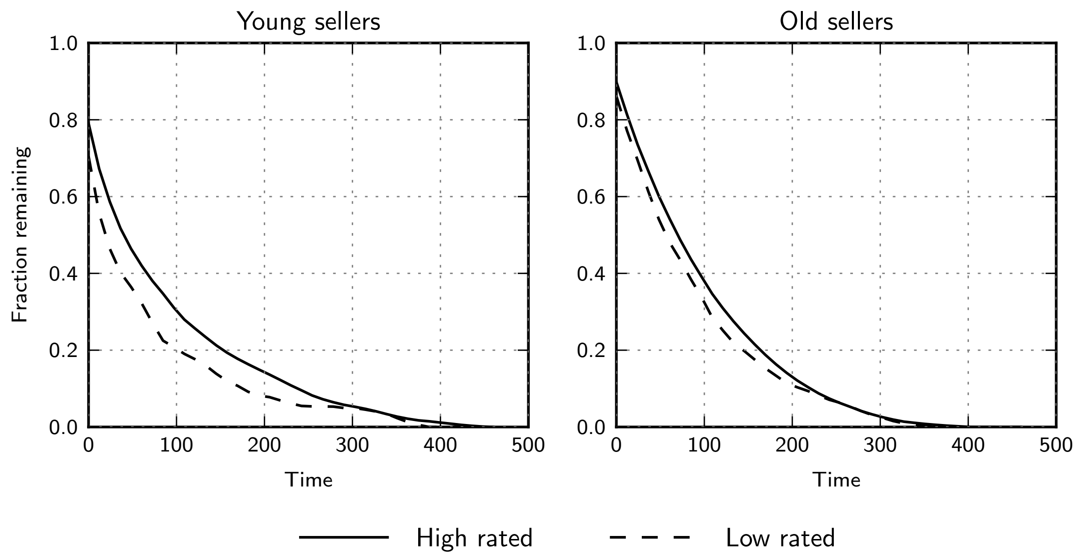

Figure 10: Survival probabilities by age and rating.

Note: Left: The fraction of young sellers with a high rating (solid) versus sellers with a low rating (dashed) remaining in the market after some time. Right: The fraction of old sellers with a high rating remaining versus sellers with a low rating remaining. High rating is defined to be higher than 4.95, low rating is defined to be less than 4.95, old is defined as greater than 50 days in the market, young is defined as less than 50 days in the market.

low-rated sellers exit the market rather than lower their price. As low quality sellers exit, the rating dispersion narrows, and the average rating increases. In the mature stage, the mature sellers continue to operate, almost all of them relatively highly rated. Sellers whose rating falls exit the market rather than face the prospect of lower prices. Thus, there is a relatively narrow spread in the rating among mature sellers.

The lack of a relationship between rating and price can be explained by a combination of two factors: First, low-rated sellers exit, leading to an omitted variable bias. If they were to stay in the market and continue to sell a product, they would sell it at a much lower price, but they exit instead. This channel is consistent with prior literature on reputation which finds strong effects of negative reviews on exit from the markets (Cabral and Hortaçsu, 2010). Second, there is a composition bias introduced by failing to control for the number of reviews left for a vendor, and the joint relationship between reviews and ratings is non-linear.

4 Model

In this section, we develop a model of a market platform that features ratings, prices, and exit. Sellers are long lived and differ in their privately observed quality. The market platform provides publicly observable information about each seller. This publicly observable information may include the seller's average rating and the seller's total number of sales made. Instead of modeling buyers and their behavior directly, we will make two simplifying assumptions.

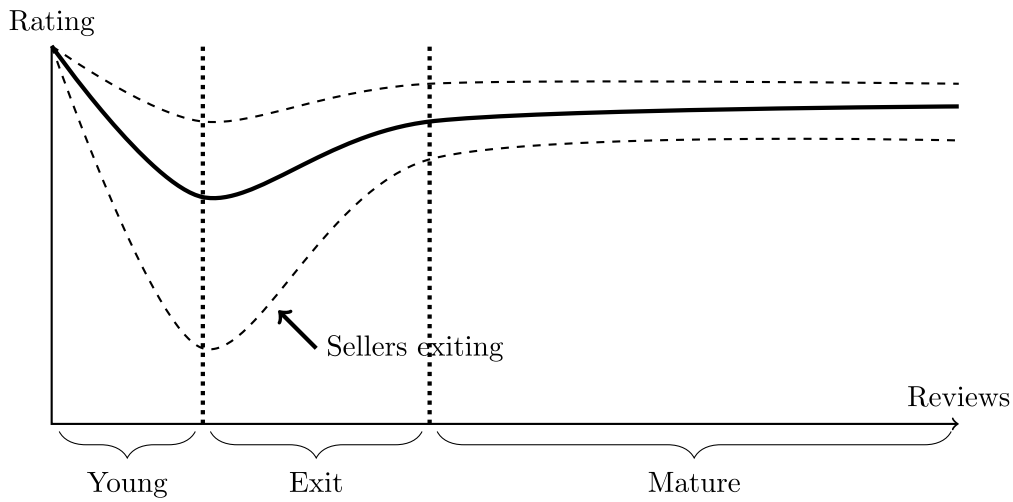

Figure 11: Graphical illustration of the life-cycle of sellers in the data.

Note: The Figure illustrates how the life-cycle of sellers interacts with ratings and the dispersion thereof. We argue that the evolution of ratings and their dispersion is mostly driven by a composition effect. Initially, all sellers (good and bad ones alike) start out at the best possible rating. As time goes by and sales are made, good sellers receive good reviews and bad sellers receive bad reviews. As bad sellers accumulate bad reviews, their rating deteriorates and they exit the market. Therefore, as sellers mature, the average rating increases and the dispersion of ratings decreases.

First, our model takes as a primitive a *rating system* (Ekmekci, 2011). The rating system specifies how a seller's publicly available information evolves over time conditional on the true quality of the seller. Second, our model takes as a primitive a demand process.[6] We assume that the price that a seller charges and the quantity sold is a function of market beliefs of the seller's quality. Market beliefs are determined in equilibrium by the rating system and seller's actions. Each period, faced with a price, a per-period cost shock, and some number of expected sales, each seller makes a decision whether or not to (irreversibly) exit. Market beliefs over a seller's quality are formed, in equilibrium, conditional on the publicly observable information about the seller that is provided by the rating system, as well as seller's exit strategies. In equilibrium, sellers optimally exit the market, taking market beliefs as given.

4.1 Environment

Time is discrete and infinite, $t = 1, 2, \ldots$. There is a continuum of sellers. Sellers are long-lived, forward looking and maximize profits, and they discount the future with discount factor $\beta \in (0, 1)$. Each seller is characterized by a privately-known and *permanent* seller quality, $\theta \in \Theta = \{\underline{\theta}, \overline{\theta}\} \subset \mathbb{R}$, satisfying $\underline{\theta} < \overline{\theta}$, and a publicly observable state $\omega \in \Omega$, which is *time-varying* and which we

[6] We take demand as exogenous because, in practice, endogenizing demand proves to be intractable. If a seller is allowed to choose a price, then a seller's choice of price should be informative about his quality. In practice this results in multiple equilibria. Saeedi (2014) takes an alternate approach and assumes sellers choose a price, but buyers do not form inferences about the quality of the seller from that price.

interpret as the seller's rating and sales made.[7][8] Seller's publicly observable state evolves according to a Markov process, conditional on θ, the seller's type. We denote the type-dependent transition probability matrix by Π_θ, a square matrix of size $|\Omega| \times |\Omega|$, whose elements represent the transition probabilities between states ω and ω':

$$\Pi_\theta(\omega, \omega') = \Pr(\omega_{t+1} = \omega' \mid \omega_t = \omega, \theta).$$

Each period, sellers incur a fixed cost of operation c. This cost is identically and independently distributed and drawn from a cumulative distribution function F. Sellers can avoid paying the fixed cost c by irreversibly exiting the market. See Figure 12 for an illustration of the sequence of events within a period.

Each period, a continuum of young sellers enters the market We denote the mass of entering sellers with private type θ and initial public state ω by $\eta_\theta(\omega)$. Under the interpretation that ω represents the publicly available information about a seller, a reasonable assumption, which we impose when estimating the model, is that there is a initial state ω_0, representing the information that a seller is new, such that

$$\eta_\theta(\omega_0) > 0, \eta_\theta(\omega) = 0 \, \forall \omega \neq \omega_0.$$

That is, all entering sellers are assigned state ω_0. For now, we allow $\eta_\theta(\omega)$ to be any measure on Ω.

We model sellers as price takers and take demand as exogenously given. Prices and demand depend only on the market beliefs about the average quality of a seller, conditional on the state. Market beliefs about the average quality of a seller at time t are denoted $\hat{\theta}_t(\omega) \in [\underline{\theta}, \overline{\theta}]$, conditional only on the publicly observable state variable. Prices and demand are then denoted $p(\hat{\theta}_t(\omega))$ and $q(\hat{\theta}_t(\omega))$ Later, for purposes of estimation, we assume that prices and quantities are given by the specific functional form

$$p(\hat{\theta}_t(\omega)) = \hat{\theta}_t(\omega), \tag{2}$$

$$q(\hat{\theta}_t(\omega)) = \gamma_0 + \gamma_1 \hat{\theta}_t(\omega), \tag{3}$$

That is, we later assume that the price simply equals the market's beliefs about the quality of the seller and that demand is a linear function of market beliefs, parameterized by β_0, β_1. This is intended to capture, in a reduced-form way, some process through which prices are determined in which higher buyer beliefs about the quality of a seller results in higher prices to that seller. Directly modeling such a process is complicated, and so, we abstract away from it.

The mass of sellers of quality θ in state ω at time t is denoted $\mu_{t\theta}(\omega)$. The age of a seller is

[7] In reality, buyers observe other characteristics of the seller, such as the five most recent reviews, and the position of the seller in the search ranking. We currently do not allow the buyers to condition on these characteristics, however, augmenting the model to allow for this is straightforward and a possible avenue for future work.

[8] Allowing the seller to choose θ, or allowing θ to be time-varying, is a potential way to introduce moral hazard into the model. Here, we focus on the problem of adverse selection, and so θ is fixed.

denoted by $a \geq 0$, and a seller's history contains the history of seller states and cost shocks,

$$h^a = (\omega_1, c_1, \omega_2, c_2, \ldots, \omega_a, c_a).$$

The set of all seller histories is $\mathcal{H}$. A seller's strategy, τ, maps time, a seller's history, and his realized cost shocks into a probability of remaining in the market, denoted $(h^a) \overset{\tau_t}{\mapsto} [0,1]$.

To represent transition probabilities in a compact way, we let $\vec{\tau}_{t\theta}, \vec{\mu}_{t\theta}, \vec{\eta}_\theta$ denote the $|\Omega|$-length vectors containing the exit probabilities, mass of sellers in each state, and mass of sellers entering into each state, that is

$$\vec{\tau}_{t\theta} = \langle \mathbb{E}_c[\tau_{t\theta}(\omega, c)] \rangle_{\omega\in\Omega}, \quad \vec{\mu}_{t\theta} = \langle \mu_{t\theta}(\omega) \rangle_{\omega\in\Omega}, \quad \text{and } \vec{\eta}_{t\theta} = \langle \eta_{t\theta}(\omega) \rangle_{\omega\in\Omega}.$$

We may then write the distribution of sellers over states as satisfying the transition rules

$$\vec{\mu}_{t+1\theta} = \Pi_\theta(\vec{\mu}_{t\theta} \circ \vec{\tau}_{t\theta}) + \vec{\eta}_\theta, \tag{4}$$

where $\circ$ denotes the pointwise product.[9]

If a seller exits, then he receives a payoff of zero for the remainder of the game. If instead a seller remains in the game, then a seller in period t, in state ω_a, with realized cost shock c_a recieves a flow payoff of

$$u_t(h^a, \hat{\theta}_t) = p(\hat{\theta}_t(\omega_a))q(\hat{\theta}_t(\omega_a)) - c_a.$$

Under the functional form assumptions (2) and (3), the flow payoff may be written

$$u_t(h^a, \hat{\theta}_t) = \hat{\theta}_t(\omega_a)(\gamma_0 + \gamma_1\hat{\theta}_t(\omega_a)) - c_a.$$

A seller's continuation value is the discounted sum of his flow payoff, until exit. The seller's entire payoff, given some exit strategy τ, entrance at time t, starting state ω, and market beliefs $\hat{\theta}$, may be written

$$U_t(\omega_0, \hat{\theta}, \tau) = \mathbb{E}\left[\sum_{a=0}^{\infty} \beta^a u_t(\tilde{h}^a, \hat{\theta}_{t+a}) \prod_{a'=0}^{a} \tau(\tilde{h}^{a'}) \mid \omega_0\right],$$

where the expectation is taken with respect to the probability measure over all histories, beginning in state ω_0, induced by Π_θ and the distribution of cost shocks; the random variable is denoted $\tilde{h}^a$. The entire payoff of a seller is the discounted sum of his flow payoffs, $u_t(\tilde{h}^a, \hat{\theta}_{t+a})$, multiplied by the probability that the seller has remained in the market up to age a, $\prod_{a'=0}^{a} \tau(\tilde{h}^{a'})$.

In the remainder of the paper, we are mainly interested in stationary equilibria. In a stationary equilibrium, market beliefs are independent of time, and seller strategies are independent of time and age. To denote stationarity, we drop the dependence on time, that is, $\hat{\theta}_t(\omega) \equiv \hat{\theta}(\omega), \tau_t(h^a) \equiv$

[9] The *pointwise product* of two vectors of length n, $\vec{x}^1, \vec{x}^2$, is the vector $\langle x_1^1 x_1^2, x_2^1 x_2^2, \ldots, x_n^1, x_n^2 \rangle$.

$\tau(\omega_a, c_a)\,\forall t, h^a$. The corresponding stationary payoffs are then denoted $u(\omega, c, \hat{\theta})$, and $U(\omega, \hat{\theta})$, functions of the current state, ω, the cost shock, c, and market beliefs, $\hat{\theta}$. The transition rules under stationarity may be written simply as

$$\vec{\mu}_\theta(\omega) = \Pi_\theta(\vec{\mu}_\theta \circ \vec{\tau}_\theta) + \vec{\eta}_\theta. \tag{5}$$

4.2 Equilibrium

Our equilibrium notion requires that sellers be exiting optimally, and that the beliefs buyers form be consistent with behavior of sellers. We focus on stationary equilibria, and so for compactness the general definition of equilibrium is omitted.

Definition An exit strategy and beliefs $\langle \tau, \hat{\theta} \rangle$ is a *stationary equilibrium* iff τ solves the seller's problem,

$$\tau(\omega, c) \in \arg\max_{\hat{\tau}} U(\omega, \hat{\theta}, \hat{\tau})\,\forall \omega \in \Omega \tag{6}$$

and buyer beliefs are consistent with Bayes rule and equilibrium seller behavior, that is, μ satisfies the stationarity condition (5), and

$$\hat{\theta}(\omega) = \underline{\theta} + (\overline{\theta} - \underline{\theta}) \frac{\mu_{\overline{\theta}}(\omega)}{\mu_{\underline{\theta}}(\omega) + \mu_{\overline{\theta}}(\omega)}, \tag{7}$$

where well-defined.[10]

We note finally that in a stationary equilibrium, in a standard result, the seller's payoff and strategy may be written recursively as a value function $V_\theta(\omega, c)$, and an exit decision, $\tau(\theta, c) \in [0, 1]$, functions of the current state, ω, and the current realization of the cost shock, c, in the standard Bellman form,

$$V_\theta(\omega, c) = \max_{\tau \in [0,1]} u(\omega, c, \hat{\theta}) + \beta \mathbb{E}[V_\theta(\tilde{\omega}, \tilde{c}) \mid \omega], \tag{8}$$

where $\tilde{\omega}, \tilde{c}$ denote the next-period draws of the state variable and the cost shock conditional on the current state.

We maintain the following assumptions for the remainder of the paper on the rating system, Π:

A1 (*Irreducible and positive recurrent rating system.*) Π_θ is irreducible (that is, for every ω, ω', there exists $n \in \mathbb{Z}$ so that $\Pi_\theta^{(n)}(\omega', \omega) > 0$), and Π is positive recurrent, i.e.

$$\sum_{n=1}^{\infty} n \cdot \Pi_\theta^{(n)}(\omega, \omega) < \infty, \forall \omega \in \Omega.$$

[10] Assumptions A1 and A2 ensure that $\mu_{\underline{\theta}}(\omega) + \mu_{\overline{\theta}}(\omega) > 0$, so that (7) is well-defined everywhere in equilibrium, so that the consistency requirement (7) holds in all states.

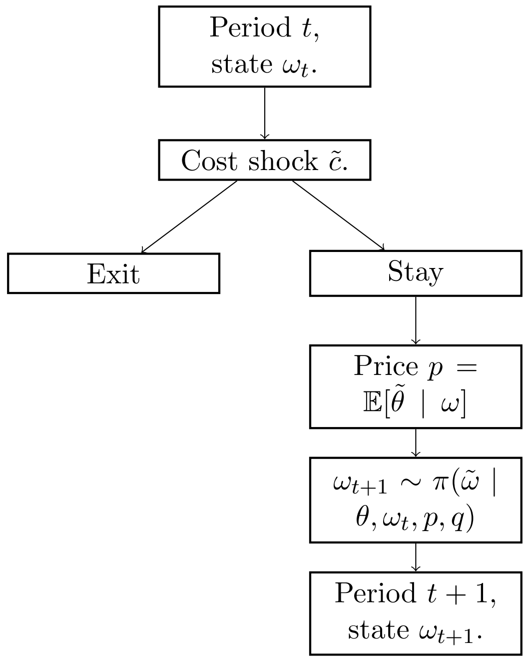

Figure 12: Sequence of events for an individual seller within one period.

A2 (*Unbounded costs.*) Seller costs satisfy

$$0 < F(\underline{\theta}) < F(\overline{\theta}) < 1,$$

and the following technical assumption,

$$\tfrac{d}{dc}\left(F(c)(c - \mathbb{E}(\tilde{c} \mid \tilde{c} \leq c))\right) > 0 \, \forall c \in \mathbb{R}. \tag{9}$$

Together, assumptions A1 and A2 ensures that there is always a positive mass of sellers in every state (possibly very small) who never exit, and therefore that beliefs over seller types conditional on a state are well-defined for all states.

The value to the seller, (8), of staying in the market is strictly decreasing in the cost shock, c, hence, in every equilibrium the exit decision is described by a cutoff strategy:

Proposition 1 (Seller's optimal exit decision)**.** *Let $\langle \tau, \hat{\theta} \rangle$ be an equilibrium. Then τ is described (up to zero-probability events) by a cutoff strategy in c,*

$$\tau_\theta(\omega, c, \hat{\theta}) = \begin{cases} 0 & c < \underline{c}_\theta(\omega, \hat{\theta}) \\ 1 & c \geq \underline{c}_\theta(\omega, \hat{\theta}) \end{cases},$$

where $\underline{c}_\theta(\omega)$ is the cutoff, satisfying

$$\underline{c}_\theta(\omega, \hat{\theta}) = \hat{\theta}(\omega) + \beta \mathbb{E}[V_\theta(\tilde{\omega}', \tilde{c}) \mid \omega]. \tag{10}$$

Proof. Immediate from the Bellman equation 8. When $c = \underline{c}_\theta(\omega, \hat{\theta})$, the seller is indifferent between staying and exiting and the seller's optimization problem has multiple solutions; by assumption F is continuous and so this is a zero probability event. □

In principle, we might compute the equilibrium value function by iterating on (8). However, it is useful to reduce the dimensionality of the problem and iterate on the expected value function instead, where we take expectation with respect to the cost shocks, $\tilde{c}$. From (10), we can write the expected value function as

$$\mathbb{E}[V_\theta(\omega, \tilde{c}, \hat{\theta})] = \underline{c}_\theta(\omega, \hat{\theta}) - \mathbb{E}[\tilde{c} \mid \tilde{c} \leq \underline{c}_\theta(\omega, \hat{\theta})]. \tag{11}$$

Together with (10), this yields

$$\underline{c}_\theta(\omega, \hat{\theta}) = \hat{\theta}(\omega) + \beta \mathbb{E}[\underline{c}_\theta(\tilde{\omega}, \hat{\theta}) - \mathbb{E}[\tilde{c} \mid \tilde{c} \leq \underline{c}_\theta(\tilde{\omega}, \hat{\theta})] \mid \theta, \omega]. \tag{12}$$

In the special case in which $\tilde{c} \sim U[0, 1]$ and $\Theta = \{0, 1\}$, (12) may be written explicitly as

$$\underline{c}_\theta(\omega) = \hat{\theta}(\omega) + \beta \sum_{\omega'} \frac{1}{2} \Pi_\theta(\omega, \omega') \underline{c}_\theta(\omega'). \tag{13}$$

A stationary equilibrium is therefore fully characterized by cutoff rules $\underline{c}$ and a distribution μ, jointly satisfying the equilibrium conditions, (6) and (7). In practice, we compute equilibria by iterating jointly on (12) and (5) until convergence is reached.

Under Assumptions A1 and A2, together with high discount rates, equilibria are unique, as summarized in the following proposition.[11]

Proposition 2. *There exists $\overline{\beta} < 1$ such that, when $\beta \geq \overline{\beta}$, the model admits a unique equilibrium.*

Proof. TBC. □

Proposition 2 establishes that, when the discount rate is sufficiently high, there is a unique equilibrium. The proof technique, which applies a global inverse function theorem (Gale and Nikaido, 1965) to establish uniqueness, also yields a method for computing $\overline{\beta}$, the limit above which uniqueness is guaranteed. Proposition 2 is a useful result empirically because generally, in models in which beliefs enter into payoffs as in (8), there are multiple equilibria, complicating their empirical analysis.

4.3 Illustration

To illustrate some dynamics of the model, we consider a special case in which the state space consists of two possible ratings, elements of $\{L, H\}$, and two possible sales levels, elements of $\{1, 2\}$, and there are two seller qualities, $\underline{\theta} = 0$ and $\overline{\theta} = 1$. There are therefore four possible states,

[11] Equilibria are not, however, guaranteed to exist.

$\{L1, H1, L2, H2\}$. A seller of quality θ with sales level s transitions between ratings according to the transition matrices

$$\Pi_\theta = \begin{pmatrix} 1-(1-\frac{s}{2})(1-\gamma_\theta) & (1-\frac{s}{2})(1-\gamma_\theta) \\ 1-(1-\frac{s}{2})(1-\gamma_\theta) & (1-\frac{s}{2})(1-\gamma_\theta) \end{pmatrix}, \tag{14}$$

where $\gamma_\theta \in (0,1)$ parameterizes the quality of the rating system. So that a high rating is associated with a high type, we assume $\gamma_0 > \frac{1}{2} > \gamma_1$. When $s = 0$, the transition matrix is simply

$$\Pi_\theta = \begin{pmatrix} \gamma_\theta & 1-\gamma_\theta \\ \gamma_\theta & 1-\gamma_\theta \end{pmatrix},$$

when $s = 1$, it is

$$\Pi_\theta = \begin{pmatrix} 1-\frac{1-\gamma_\theta}{2} & \frac{1-\gamma_\theta}{2} \\ \frac{\gamma_\theta}{2} & 1-\frac{1-\gamma_\theta}{2} \end{pmatrix}.$$

With some exogenous probability $\rho \in (0,1)$, sellers have a chance of transitioning from the low sales rate, $s = 0$, to the high sales rate, $s = 1$. Once a seller reaches a high sales rate, they cannot transition back to a low sales rate. This parameterization of the model is one in which low quality sellers are more likely than the high quality sellers to transition to the low rating, but when $s = 1$, they are less likely to transition between ratings, capturing, in a reduced form way, the idea that a seller who has already made many sales needs more bad reviews to shift their average rating.

What is the equilibrium? Let $\underline{c}_{\omega_r\omega_s\theta}$ represent the exit cutoff for type θ in state (ω_r, ω_s), and let $\mu_{\omega_r\omega_s\theta}$ denote the measure of type θ in state (ω_r, ω_s). The stationarity conditions may be written explicitly as

$$\overbrace{\mu_{H1\theta}\left((1-c_{H1\theta}) + c_{H1\theta}\left(\frac{1-\gamma_{1-\theta}}{2} + \rho\right)\right)}^{\text{Mass exiting}} = \overbrace{\tfrac{1}{2} + \mu_{L1\theta}c_{L1\theta}\left(\frac{1-\gamma_\theta}{2}\right)}^{\text{Mass entering}} \tag{$H1\theta$}$$

$$\mu_{L1\theta}\left((1-c_{L1\theta}) + c_{L1\theta}\left(\frac{1-\gamma_\theta}{2} + \rho\right)\right) = \mu_{H1\theta}c_{H1\theta}\left(\frac{1-\gamma_{1-\theta}}{2}\right) \tag{$L1\theta$}$$

$$\mu_{H2\theta}\left((1-c_{H2\theta}) + c_{H2\theta}\left(\frac{1-\gamma_{1-\theta}}{4}\right)\right) = \mu_{H1\theta}c_{H1\theta}\rho + \mu_{L2\theta}c_{L2\theta}\left(\frac{1-\gamma_\theta}{4}\right) \tag{$H2\theta$}$$

$$\mu_{L2\theta}\left((1-c_{L2\theta}) + c_{L2\theta}\left(\frac{1-\gamma_\theta}{4}\right)\right) = \mu_{L1\theta}c_{L1\theta}\rho + \mu_{H2\theta}c_{H2\theta}\left(\frac{1-\gamma_{1-\theta}}{4}\right) \tag{$L2\theta$}$$

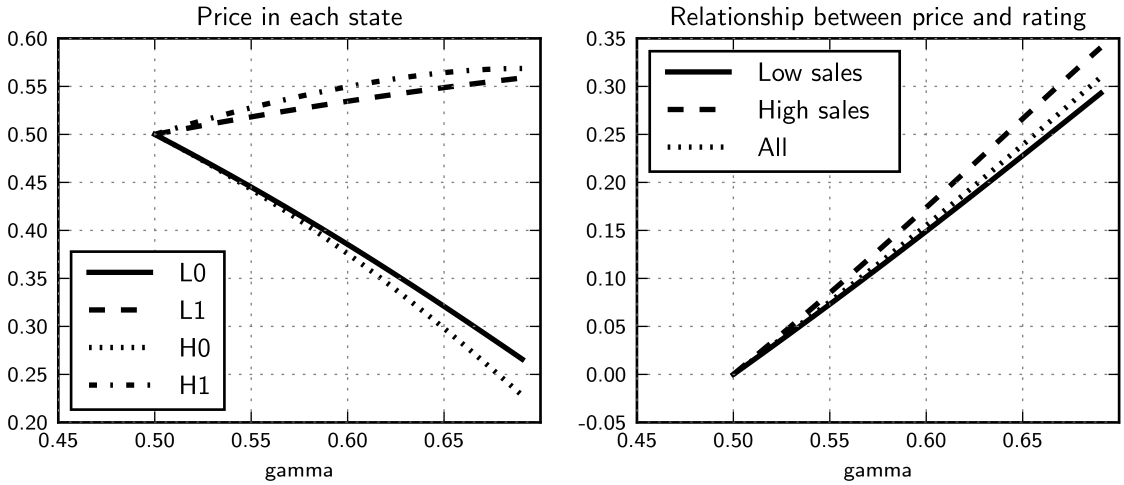

Figure 13: The price in each state, as function of the precision of the rating system, γ.

Note: Left: The price, as a function of each of the possible four states, solid ($L1$, low rating, low sales), dashed ($H1$, high rating, low sales), dotted ($L2$, low rating, high sales), and dot-dashed ($H2$, high rating, high sales). As the rating system becomes more accurate, the prices received by low rated and highly-rated sellers diverges. Right: The relationship between rating and price (dotted) is increasing in the precision of the rating system, γ. Controlling for sales reveals evidence of a composition effect—sellers with higher sales see a larger relationship between price and rating.

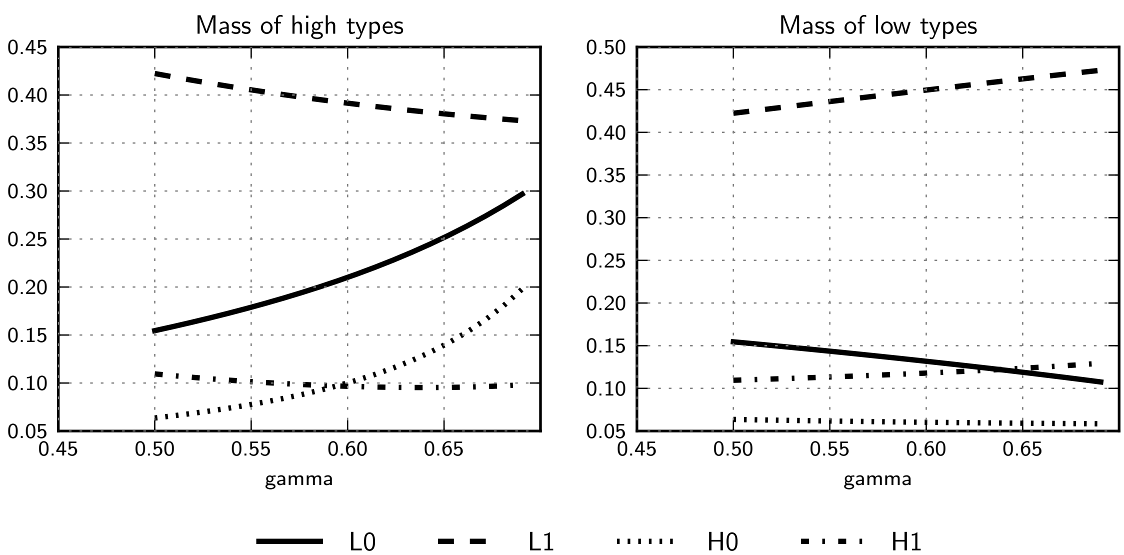

Figure 14: Mass of sellers in each state, low rated, low sales (solid), high rated, low sales (dashed), low rated, high sales (dotted), and high rated, high sales (dash-dotted).

Note: The figures show how the mass of types varies (due to exit) as a function of the precision of the rating system, γ. The left figure shows the mass of high types, the right figure the mass of low types.

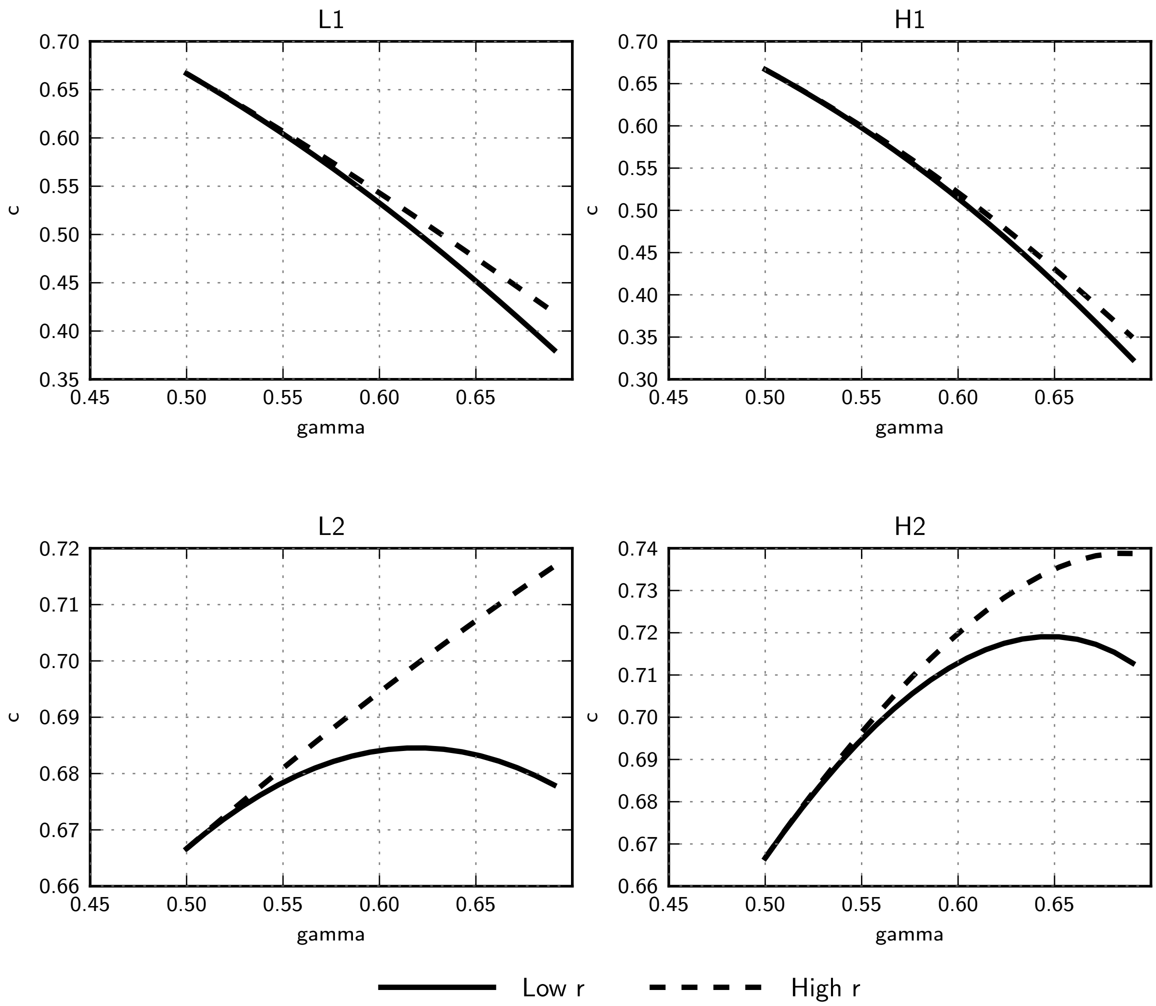

Figure 15: The cost shock cutoff in each of the four possible states, for low quality sellers (solid), and high quality sellers (dashed), as a function of the quality of the rating system.

Note: The figures show the cost shock cutoff $\underline{c}$ as a function of the precision of the rating system, γ. Solid lines refer to high quality sellers, dashed lines to low quality sellers. Each figure represents each of the four possible states, $L1$ (low rating, low sales), $H1$ (high rating, low sales), $L2$ (low rating, high sales), and $H2$ (high rating, high sales).

The corresponding optimality conditions, from (13), are

$$\underline{c}_{H1\theta} = \frac{\mu_{H11}}{\mu_{H11} + \mu_{H10}} + \frac{\beta}{2}\left(p\underline{c}_{H2\theta} + \left(\frac{1-\gamma_{1-\theta}}{2}\right)\underline{c}_{L1\theta} + \left(1 - p - \frac{1-\gamma_{1-\theta}}{2}\right)\underline{c}_{H1\theta}\right) \qquad (cH1\theta)$$

$$\underline{c}_{L1\theta} = \frac{\mu_{L11}}{\mu_{L11} + \mu_{L10}} + \frac{\beta}{2}\left(p\underline{c}_{L2\theta} + \left(\frac{1-\gamma_{\theta}}{2}\right)\underline{c}_{H1\theta} + \left(1 - p - \frac{1-\gamma_{\theta}}{2}\right)\underline{c}_{L1\theta}\right) \qquad (cL1\theta)$$

$$\underline{c}_{H2\theta} = \frac{\mu_{H21}}{\mu_{H21} + \mu_{H20}} + \frac{\beta}{2}\left(\left(\frac{1-\gamma_{1-\theta}}{4}\right)\underline{c}_{L2\theta} + \left(1 - \frac{1-\gamma_{1-\theta}}{4}\right)\underline{c}_{H2\theta}\right) \qquad (cH2\theta)$$

$$\underline{c}_{L2\theta} = \frac{\mu_{L21}}{\mu_{L21} + \mu_{L20}} + \frac{\beta}{2}\left(\left(\frac{1-\gamma_{\theta}}{4}\right)\underline{c}_{H2\theta} + \left(1 - \frac{1-\gamma_{\theta}}{4}\right)\underline{c}_{L2\theta}\right). \qquad (cL2\theta)$$

This is a system of eight equations in eight unknowns, and may be solved explicitly under the assumption that $\gamma_0 = 1 - \gamma_1 = \gamma$. (The algebra is omitted.) Figures 13 and 15 display the prices, relationship between price and rating, and exit probabilities, respectively. We highlight the following relevant features of the simple model:

1. The relationship between the price and the rating is dependent on the number of sales made by the seller (Figure 13).

2. The spread between the price in a high-rated state and a low-rated state is dependent on the quality of the rating system, γ (Figure 13).

3. High quality sellers are more likely to exit compared to low quality sellers, but the exit decision is non-monotonic in the quality of the rating system (Figure 15).

4. The mass of low quality sellers who are low ranked is increasing in the quality of the rating system, and the mass of high quality sellers who are high ranked is decreasing in the quality of the rating system (Figure 14).

5 Estimation

In this section, we discuss how we bring the model to the data. When we estimate the model, we only consider the sale of Cannabis, the product category for which we have the most observations. This leaves us with a total of 279,054 reviews across 824 vendors (for summary statistics see Table 1). The restriction to one product category simplifies the analysis, albeit at the cost of ignoring reputation effects that span product categories. A vendor who accumulates a reputation selling one product category, may benefit from this reputation when he sells a different product category. However, for 730 of the 824 vendors in our data, Cannabis is the only product sold during the seller's lifetime.

As in Section 3, we use reviews left as a measure of sales made. We aggregate our data into a weekly panel data set of vendors. We normalize all prices by weight, which allows us to treat multiple listings of the same vendor as the same product. We further normalize prices by removing a shipping location fixed effect. This makes prices comparable across geographic regions, which we do not model explicitly. For vendors with multiple sales per week, we use the end-of-week realizations of the number of sales and rating and the week's average realization of the price.

We estimate the model with maximum likelihood using a nested fixed point algorithm as in Rust (1987). For the purpose of estimating the model, the cost shocks serve as the model's econometric error. These shocks are observed by the vendor but unobserved by the econometrician. A period in the model corresponds to a week. For each vendor we observe how the vendor's ratings and sales evolve over time. We also observe when a vendor decides to exit. Importantly, we do not observe vendor types $\theta \in \{\underline{\theta}, \overline{\theta}\}$. These are private information. We model these types as finite mixtures as commonly done in single-agent dynamic discrete choice problems (e.g. in Eckstein and Wolpin (1990, 1999) or Keane and Wolpin (1997)).

5.1 Model parameterization

To estimate the model, we impose some additional structure on parts of it. In particular, we parameterize the rating system and the demand process. Each seller's publicly observable information, $\omega \equiv [r, s]$, consists of a rating, r, and the total number of sales made, s. Ratings can take values in $\{0.00, 0.01, 0.02, \ldots, \ldots, 4.99, 5.00\}$, while sales can take values in $\{0-1, 1-5, 5-10, \ldots, 5000-\infty\}$. We model the rating of seller i at age a, $\tilde{r}_{ia}$, as evolving stochastically according to a Markov process dependent only on the total number of sales made, the seller's type, and the previous period's rating,

$$\tilde{r}_{ia+1} = \xi r_{ia} + (1-\xi)\rho_\theta + \varepsilon_{ia} \quad \text{with} \quad \varepsilon_{ia} \sim \mathrm{N}(0, \sigma_r), \tag{15}$$

which is governed by the parameters, ξ, $\langle \rho_\theta \rangle_{\theta \in \{\underline{\theta}, \overline{\theta}\}}$ and σ_r. The parameter ρ_θ measures the 'true' rating of a type to which a type θ-seller's rating converges in the long-run. The coefficient $\xi \in (0, 1)$ governs how slowly the rating adjusts. When $\xi = 1$, the rating system is independent of a seller's type, and so is uninformative about type. When $\xi = 0$, the rating of a seller is, on average, ρ_θ. When $\xi = 0$ and σ_r is small, the rating almost entirely reveals the seller's type. $\sigma_r > 0$ describes the standard deviation of the ratings process.[12] The process is then discretized using the technique of Tauchen (1986).

Sales are assumed to evolve exogenously and independently of a seller's type and rating. Specifically, with probability γ, sellers graduate from one sales level to the next.[13] We denote the resulting

[12]We intend for this particular specification to capture the true underlying stochastic process, in which buyers purchase the product, sample it, and leave reviews, which are then averaged to form the rating. In an alternate specification, we model this stochastic process directly, however, this approach introduces computational difficulties, as well as some assumptions about the underlying process through which buyers leave reviews. To avoid this difficulty, we abstract away and use the specification (15). A more general specification of (15) may allow the variance of the process to be time-varying, to capture the concept that as more reviews are left, each additional review will have less of an impact on the rating. On the other hand, as more reviews are left, sellers tend to receive more reviews, increasing the variance of the ratings process, and so it is not clear *ex-ante* whether the variance of the ratings process should increase or decrease as more sales are made. In an alternate specification of the model, we allow the variance of the rating process to be $\sigma_r \lambda^{s_{ia}}$, where $\lambda > 0$ is a parameter capturing a time-varying variance of the ratings process, but in practice, we estimate λ very close to 1, and so omit it here for simplicity.

[13]More generally, we might wish to allow the sales made to depend on market beliefs about the seller. This

transition probability matrices that characterize the joint evolution of publicly observable information ω by $\pi_{\overline{\theta}}(\omega, \omega')$ and $\pi_{\underline{\theta}}(\omega, \omega')$. We parameterize the initial distribution of publicly observable information for new sellers by $\eta(\omega)$, which we assume does not depend on the seller's quality θ. Costs shocks are assumed to be normally distributed with mean μ_c and standard deviation σ_c. The distribution of cost shocks is assumed to be independent of the seller's type, θ.

In the model, all vendors with the same publicly observable state, ω, receive the same price. In the data, this is clearly not true. To ensure that the model is statistically non-degenerate, we add identically and independently distributed measurement error to the price. We assume that the measurement error follows a log-normal distribution with the location parameter fixed at zero and scale σ_p. The log-normal specification both excludes non-positive prices, and as well, fits the true distribution of prices (conditional on state) well. This measurement error is independent of the seller's quality θ. We denote the density of the measurement error by ϕ.

5.2 Likelihood

We now construct the likelihood function. We observe each vendor i for A periods. A is not the same for all vendors, because different vendors exit the market at different ages. For each vendor we observe a sequence of the publicly observable state variable, $\omega_{i1}, \ldots, \omega_{iA}$, a sequence of prices, $p_{i1}, \ldots, p_{iA}$, and whether the vendor exits in the last period d_{iA}. Because of the sudden closure of the market place, we we do not observe the exit decision for all vendors.

First, consider a vendor that we see exiting the market at time A, i.e. $d_{iA} = 1$. The likelihood contribution for this vendor conditional on θ is given by

$$
\begin{aligned}
p(\omega_{i1}, \ldots, &\omega_{iA}, d_{iA} = 1|\theta) = \\
&\eta(\omega_{i1}) \Pr(\tilde{c} < \underline{c}_\theta(\omega_{i1}))\phi(p_{i1} - \mathbb{E}_\mu[\theta|\omega_{i1}]) && \textit{initial period} \\
&\prod_{a=2}^{A-1} \left[\pi_\theta(\omega_{ia-1}, \omega_{ia}) \Pr(\tilde{c} < \underline{c}_\theta(\omega_{ia}))\right] \phi(p_{ia} - \mathbb{E}_\mu[\theta|\omega_{ia}]) && \textit{intervening periods} \\
&\pi_\theta(\omega_{iA-1}, \omega_{iA}) \Pr(\tilde{c} \geq \underline{c}_\theta(\omega_{iA})) && \textit{exit period}
\end{aligned}
$$

Next, consider a vendor that we observe up until time T without exiting the market ($d_{iA} = 0$).

specification of demand fits the data well. The model when estimated with the full demand curve which depends on market beliefs generally fails to find evidence of any relationship between buyer beliefs and total sales. Since price is assumed to already capture the expected value of a seller to a buyer, it may be that any variations in sales due to buyer beliefs are of second-order importance. Regardless, even if more favorable beliefs were to lead to higher sales, the effect on reputation is then of third-order importance in the model, and this is the effect we are interested in.

The likelihood contribution for such a vendor equals

$$\begin{aligned} p(\omega_{i1},\ldots,&\omega_{iA}, d_{iA}=0|\theta) = \\ &\eta(\omega_{i1}) \Pr(\tilde{c} < \underline{c}_\theta(\omega_{i1}))) \phi(p_{i1} - \mathbb{E}_\mu[\theta|\omega_{i1}]) && \textit{initial period} \\ &\prod_{a=2}^{A} \left[\pi_\theta(\omega_{ia-1},\omega_{ia}) \Pr(\tilde{c} < \underline{c}_\theta(\omega_{ia}))\right] \phi(p_{ia} - \mathbb{E}_\mu[\theta|\omega_{ia}]) && \textit{all other periods} \end{aligned}$$

The entire log-likelihood contribution for a vendor is therefore given by

$$\ell(\omega_{i1},\ldots,\omega_{iA}, d_{iA}) = \log\left[\alpha p(\omega_{i1},\ldots,\omega_{iA}, d_{iA}|\overline{\theta}) + (1-\alpha) p(\omega_{i1},\ldots,\omega_{iA}, d_{iA}|\underline{\theta})\right],$$

where we integrate out the vendor's type $\theta \in \{\overline{\theta}, \underline{\theta}\}$ using the weights α and $1-\alpha$. The log-likelihood function is then given by

$$\sum_{i=1}^{N} \ell(\omega_{i1},\ldots,\omega_{iA}, d_{iA}).$$

We do not face the traditional initial conditions problem in the dynamic discrete choice literature, because we assumed that the initial distribution of the publicly observable information ω_{i1} is independent of the private vendor type. Also we only use vendors in the estimation for which we observe the initial sale in the data. We thereby bypass issues resulting from left-censoring.

5.3 Identification

In our model, firms' only decision is when to leave the market platform. Therefore, our model is essentially an optimal stopping problem with permanent unobserved heterogeneity. Optimal stopping problems and their identification properties have been extensively studied in the literature (see Abbring (2010) for a review). Obtaining identification in this class of models is notoriously difficult when allowing for dynamic selection on unobservables (as we do in our case). Our model differs from the traditional optimal stopping problem in the dynamic discrete choice literature in several important aspects, which allows us to obtain identification for our model. For one, we directly observe the evolution of the state variable, which we assume to be type-dependent. Following Kasahara and Shimotsu (2009), this allows us to identify the type probabilities. For another, we directly observe prices and a proxy for sales, which means that we observe vendors' flow payoffs. Ideally, this data would identify the distribution of cost shocks, parameterized by μ_c and σ_c. However, as in most dynamic discrete choice models, we cannot separately identify the location and scale of this distribution. We therefore fix the scale σ_c at 1, and only estimate the location μ_c.

5.4 Point estimates

We estimate ten model parameters. The type-dependent parameters of the model are $\langle \theta, \rho_\theta \rangle_{\theta \in \{\underline{\theta}, \overline{\theta}\}}$. The type-independent parameters are $\langle \alpha, \mu_c, \gamma, \xi, \sigma_r, \sigma_p \rangle$. We do not attempt to estimate the discount factor β, which we fix at 0.996, corresponding to an annual interest rate of approximately 25%.[14] For ease of interpretation, the model is estimated on data so that seller prices lie on the interval $[0, 1]$. In addition, we removed location effects from the price data.

We proceed in two steps. First, we obtain starting values for a subset of parameters directly from the data without solving the model. Second, using these starting values we estimate the entire model using a nested-fixed point algorithm.

We obtain starting values (or *offline*) estimates for the value of high and low quality sellers ($\underline{\theta}$ and $\overline{\theta}$), the fraction of entering high quality sellers (α), the parameter governing the growth in sales (γ), and the standard deviations of the ratings and price processes (σ_r and σ_p). The value of high quality sellers, $\overline{\theta}$, is estimated as the average price charged by high rated, mature sellers, who charge the highest prices in the market. The value of low quality sellers, $\underline{\theta}$, is estimated by looking at the average price charged by low rated, mature sellers, who charge the lowest prices in the market. The fraction of entering types who are the high type is then estimated using the formula

$$\alpha^{\text{offline}} = \frac{\mathbb{E}[p_{ia} \mid r_{ia} = 1, s_{ia} = 1] - \underline{\theta}^{\text{offline}}}{\overline{\theta}^{\text{offline}} - \underline{\theta}^{\text{offline}}}, \tag{16}$$

where $\mathbb{E}[p_{ia} \mid r_{ia} = 1, s_{ia} = 1]$ is the empirical mean of the price charged by sellers entering the market. By assumption, α is exactly the fraction of these sellers who are the high type, from which equation (16) is derived.[15] Estimates for the sales growth rate, γ^{offline}, are derived from the empirical analog

$$\gamma^{\text{offline}} = \Pr[s_{ia+1} = s_{ia} + 1 | s_{ia}],$$

and estimates for the standard deviation of the price and ratings process are derived from the empirical standard deviations

$$\sigma_r^{\text{offline}} = \sqrt{\text{var}(r_{ia+1} - r_{ia})}$$
$$\sigma_p^{\text{offline}} = \sqrt{\text{var}(p_{ia})}.$$

This produces offline estimates $\langle \underline{\theta}^{\text{offline}}, \overline{\theta}^{\text{offline}}, \alpha^{\text{offline}}, \gamma^{\text{offline}}, \sigma_r^{\text{offline}}, \sigma_p^{\text{offline}} \rangle$, which we summarize in Table 3.

[14] One period in the model corresponds to a day. Little research has been done to ascertain the discount factor of drug dealers in online marketplaces. We assume that they are impatient for a variety of reasons. They presumably face impatient suppliers, have little access to formal credit, and furthermore face risks associated with law enforcement and the sudden shutdown of the marketplace.

[15] The average price charged by new entrants equals $\alpha\overline{\theta} + (1 - \alpha)\underline{\theta}$. Solving this expression for α gives us equation (16).

Parameter	Variable	Estimate
Low type value	$\underline{\theta}^{\text{offline}}$	0.354
High type value	$\overline{\theta}^{\text{offline}}$	0.439
Fraction of entering high type	α^{offline}	0.263
Probability of moving up in sales	γ^{offline}	0.370
Standard deviation of rating process	$\sigma_r^{\text{offline}}$	0.199
Standard deviation of prices	$\sigma_p^{\text{offline}}$	0.129

Table 3: Offline parameter estimates.

With these starting values in hand, we then estimate the entire model using a nested-fixed point algorithm. For the subset of parameters for which we obtained offline estimates, we use these as starting values. For the remaining parameters we use an ad-hoc procedure to obtain starting values. Since we estimate all parameters of the model simultaneously (it corresponds to the third step in Rust's (1987) estimation procedure), we can compute asymptotic standard errors using the outer-product of the gradients to estimate the asymptotic covariance matrix. The model estimates and standard errors are reported in Table 4.

The share of entering high types, α, is estimated at 0.233. This estimate indicates that the market may indeed be prone to adverse selection, because there is a non-trivial share of high and low types. The estimates of $\underline{\theta}$ and $\overline{\theta}$ indicate that the quality difference between low types and high types is substantial. The low type value, $\underline{\theta}$, is estimated at 0.300 and the high type value, $\overline{\theta}$, is estimated at 0.525. Since we normalized prices to take on values on the unit interval, these estimates correspond to 10.5 dollars per gram and 18.4 dollars per gram, respectively.

In Figure 16 we simulate data using the estimated model and compare the moments of the model to the moments of the data. Overall, the model fits the data very well.

Parameter	Variable	Estimate	Std. error
Low type value	$\underline{\theta}$	0.300	0.001
High type value	$\overline{\theta}$	0.525	0.004
Fraction of entering high type	α	0.233	0.006
Mean of cost shocks	μ_c	0.386	0.001
Probability of making sale	γ	0.293	0.002
Average low type rating	$\rho_{\underline{\theta}}$	5.010	0.007
Average high type rating	$\rho_{\overline{\theta}}$	6.372	0.035
Rate of rating adjustment	ξ	0.060	0.001
Standard deviation of rating process	σ_r	0.037	0.001
Standard deviation of prices	σ_p	0.144	0.001

Table 4: Parameter estimates.

In Figure 17 we plot the equilibrium price and exit probabilities, at the estimated parameters, as a function of the rating, the number of sales made, and the quality of a seller. The top panel shows prices. Prices are highest for sellers with high rating and large number of sales and lowest for sellers with low rating and a large number of sales. For seller with few sales, prices are approximately independent of the seller's rating. In the bottom two panels, we show the probability of staying in the market separately for high and low quality sellers. Overall, high quality sellers are much more likely to stay in the market than low quality sellers. High quality sellers are most likely to leave when their rating is low and they have made a moderate number of sales (approximately between 1,000 and 2,000 sales). A low quality seller's probability of leaving the market is increasing in sales made when the rating is low and decreasing in sales when the rating is high.

6 Results

TBC

7 Conclusion

This paper describes the collection of a large and novel dataset of price and reputation in a black market. We replicated prior reduced form work on the relationship between price, exit, and reputation, with an emphasis on the dynamic incentives faced by firms in the market. We showed that the relationship between price and ratings is positive and increasing in the number of sales made, consistent with a model of adverse selection and consumers forming rational inferences about the quality of the product jointly from the rating and the number of sales made.

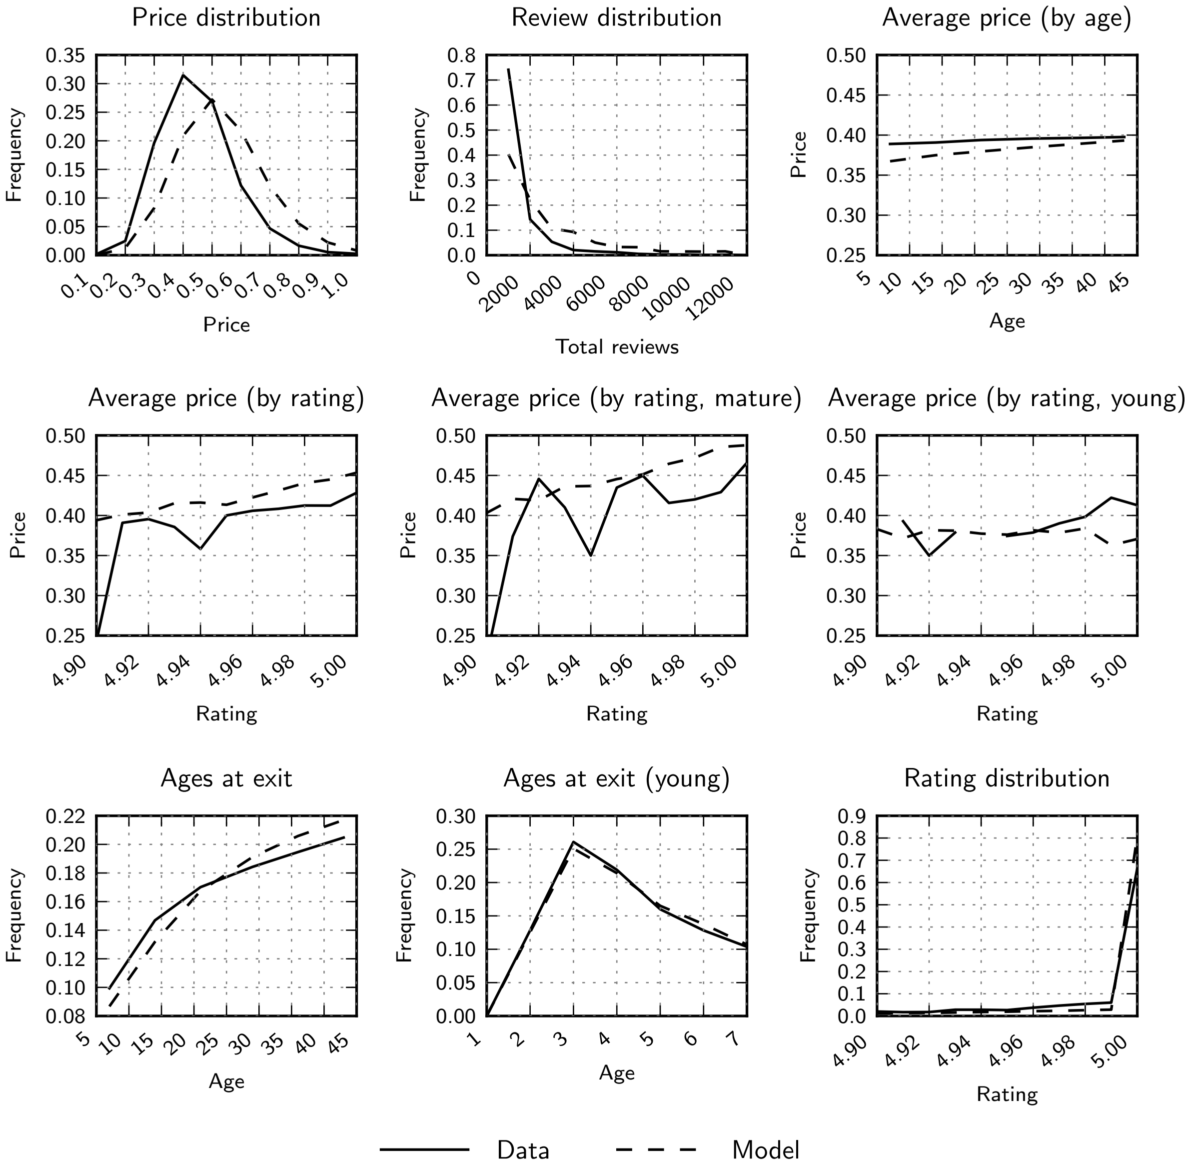

Figure 16: Model fit.

Note: Model fit graphs produced by simulation. Mature sellers are defined as those for whom there are more than 1000 reviews. Young sellers are those for whom there are fewer than 300 reviews. For ages at exit among young sellers (bottom, center), the age at exit among sellers who had been in the market fewer than 50 days is plotted. Omitted in the plots of age at exit (bottom left, bottom center) are sellers who exited the market when it closed.

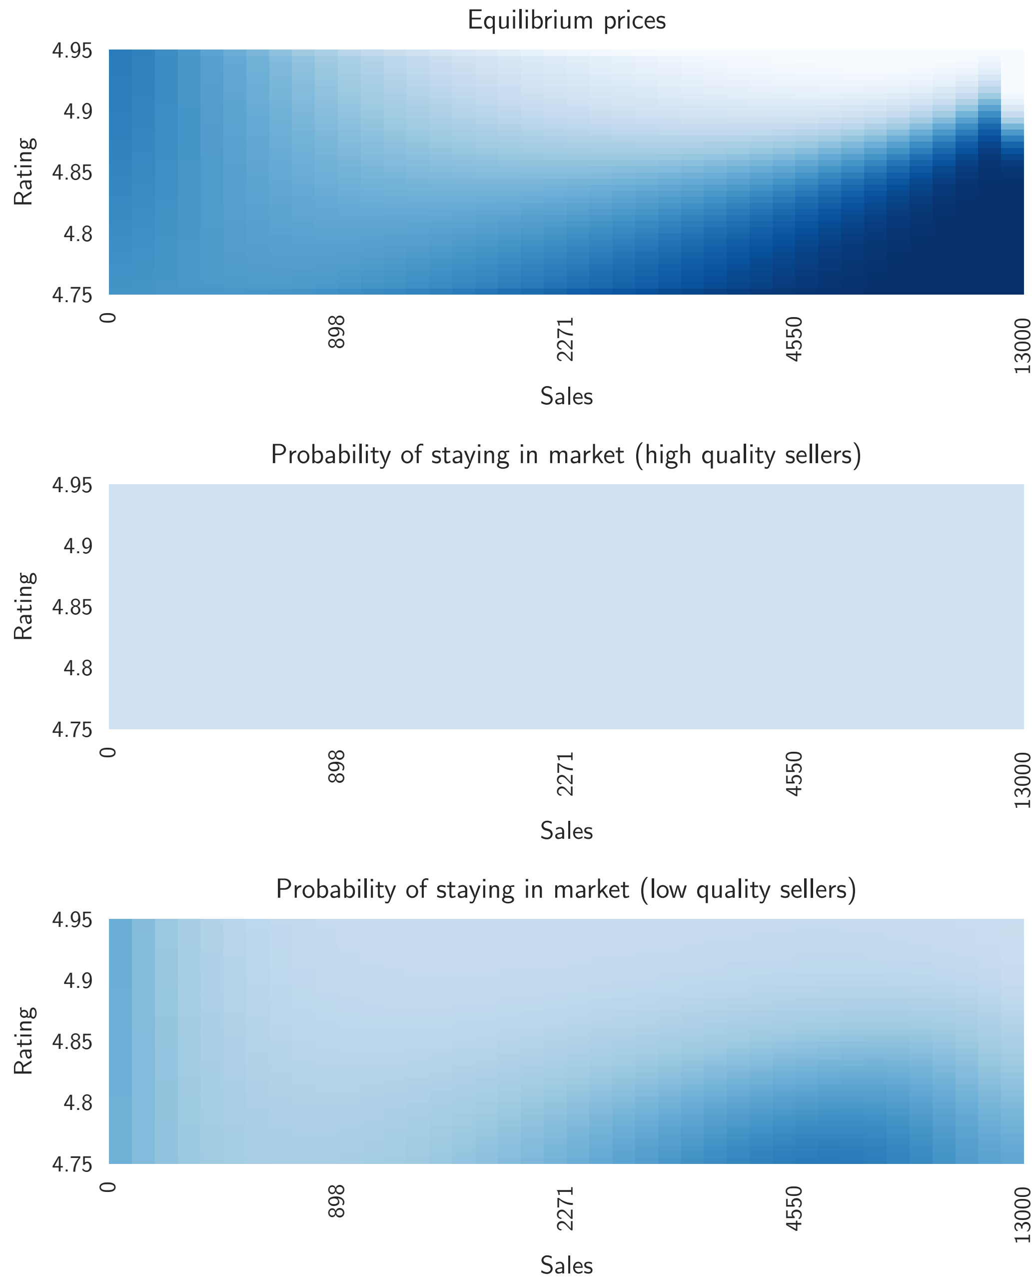

Figure 17: Prices and exit probabilities, in equilibrium, at estimated model parameters.

Note: Top: Price received in equilibrium, as a function of the rating (y-axis) and the sales (x-axis). Lighter is higher. Middle: The probability of remaining in the market, as a function of the rating and sales, for high quality sellers. Bottom: The probability of remaining in the market, as a function of the rating and sales, for low quality sellers. High quality sellers are more likely to stay in the market than low quality sellers. The gradient between price and rating is increasing in the number of sales made.

References

Abbring, J. H. (2010): "Identification of dynamic discrete choice models," *Annu. Rev. Econ.*, 2, 367–394.

Ba, S. and P. A. Pavlou (2002): "Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior," *MIS Quarterly*, 243–268.

Bajari, P. and A. Hortacsu (2003): "The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions," *RAND Journal of Economics*, 329–355.

Bohren, A. (2012): "Stochastic games in continuous time: Persistent actions in long-run relationships," *PIER Working Paper*.

Bolton, G. E., E. Katok, and A. Ockenfels (2004): "How effective are electronic reputation mechanisms? An experimental investigation," *Management Science*, 50, 1587–1602.

Cabral, L. (2012): "Reputation on the Internet," *The Oxford Handbook of the Digital Economy*, 343–354.

Cabral, L. and A. Hortaçsu (2010): "The dynamics of seller reputation: Evidence from eBay," *Journal of Industrial Economics*, 58, 54–78.

DEA (2016): "STRIDE dataset," http://www.dea.gov/resource-center/stride-data.shtml., full dataset obtained via Freedom of Information Act Request, March 14th, 2016.

Dellarocas, C. (2001): "Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms," in *Proceedings of the 3rd ACM Conference on Electronic Commerce*, ACM, 171–179.

——— (2002): "Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality," in *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, Springer, 238–252.

Dobkin, C. and N. Nicosia (2009): "The war on drugs: methamphetamine, public health, and crime," *American Economic Review*, 99, 324–349.

Douceur, J. R. (2002): "The Sybil attack," in *Peer-to-peer Systems*, Springer, 251–260.

Eaton, D. H. (2005a): "Reputation effects in online auction markets," *Available at SSRN 739765.*

——— (2005b): "Valuing information: Evidence from guitar auctions on eBay," *Journal of Applied Economics & Policy*, 24, 1.

Eckstein, Z. and K. I. Wolpin (1990): "Estimating a market equilibrium search model from panel data on individuals," *Econometrica: Journal of the Econometric Society*, 783–808.

——— (1999): "Why youths drop out of high school: The impact of preferences, opportunities, and abilities," *Econometrica*, 67, 1295–1339.

Ederington, L. H. and M. Dewally (2002): "A comparison of reputation, certification, warranties, and information disclosure as remedies for information asymmetries: Lessons from the on-line comic book market," *SSRN Working Paper Series*.

Ekmekci, M. (2011): "Sustainable reputations with rating systems," *Journal of Economic Theory*, 146, 479–503.

Fan, Y., J. Ju, and M. Xiao (2016): "Reputation premium and reputation management: Evidence from the largest e-commerce platform in china," *International Journal of Industrial Organization*, 46, 63–76.

Gale, D. and H. Nikaido (1965): "The Jacobian matrix and global univalence of mappings," *Mathematische Annalen*, 159, 81–93.

Galenianos, M. and A. Gavazza (2015): "A quantitative analysis of the retail market for illicit drugs," *CEPR Discussion Paper No. DP10363*.

Galenianos, M., R. L. Pacula, and N. Persico (2012): "A search-theoretic model of the retail market for illicit drugs," *Review of Economic Studies*, 79, 1239–1269.

Ghose, A., P. G. Ipeirotis, and A. Sundararajan (2007): "Opinion mining using econometrics: A case study on reputation systems," in *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics*.

Houser, D. and J. Wooders (2006): "Reputation in auctions: Theory, and evidence from eBay," *Journal of Economics & Management Strategy*, 15, 353–369.

Hu, J. (2015): "Biased learning and permanent reputation," *PIER Working Paper*.

Jolivet, G., B. Jullien, and F. Postel-Vinay (2016): "Reputation and prices on the e-market: Evidence from a major french platform," *International Journal of Industrial Organization*, 45, 59–75.

Kasahara, H. and K. Shimotsu (2009): "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 77, 135–175.

Keane, M. and K. Wolpin (1997): "The career decisions of young men," *Journal of Political Economy*, 105, 473–522.

Lee, Z., I. Im, and S. J. Lee (2000): "The effect of negative buyer feedback on prices in Internet auction markets," in *Proceedings of the twenty first international conference on information systems*, Association for Information Systems, 286–287.

LEVITT, S. D. AND S. A. VENKATESH (2000): "An economic analysis of a drug-selling gang's finances," *Quarterly Journal of Economics*, 115, 755–789.

MAILATH, G. AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

MCDONALD, C. G. AND V. C. SLAWSON (2002): "Reputation in an internet auction market," *Economic inquiry*, 40, 633–650.

MELNIK, M. AND J. ALM (2002): "Does a seller's ecommerce reputation matter? Evidence from eBay auctions," *Journal of Industrial Economics*, 50, 337–349.

NEWBERRY, P. AND X. ZHOU (2016): "The heterogenous effects of online reputation for local and national retailers," Tech. rep., mimeo, Penn State University.

RESNICK, P. AND R. ZECKHAUSER (2002): "Trust among strangers in internet transactions: Empirical analysis of eBays reputation system," *The Economics of the Internet and E-commerce*, 11, 23–25.

RESNICK, P., R. ZECKHAUSER, J. SWANSON, AND K. LOCKWOOD (2006): "The value of reputation on eBay: A controlled experiment," *Experimental Economics*, 9, 79–101.

RESNICK, P. ET AL. (2001): "The social cost of cheap pseudonyms," *Journal of Economics & Management Strategy*, 10, 173–199.

RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 999–1033.

SAEEDI, M. (2014): "Reputation and adverse selection, theory and evidence from eBay," *Working Paper*.

SOSKA, K. AND N. CHRISTIN (2015): "Measuring the longitudinal evolution of the online anonymous marketplace ecosystem," in *Proceedings of the 24th USENIX Security Symposium*.

TAUCHEN, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20, 177–181.

YOGANARASIMHAN, H. (2013): "The value of reputation in an online freelance marketplace," *Marketing Science*, 32, 860–891.

ZHANG, J. (2006): "The roles of players and reputation: evidence from eBay online auctions," *Decision Support Systems*, 42, 1800–1818.