The Value of Reputation in an Online Freelance Marketplace

Hema Yoganarasimhan [*]

University of California, Davis

August 6, 2013

Abstract

Online freelance marketplaces are websites that match buyers of electronically deliverable services with freelancers. While freelancing has grown in recent years, it faces the classic 'information asymmetry' problem – buyers face uncertainty over seller quality. Typically, these markets use reputation systems to alleviate this issue, but the effectiveness of these systems is open to debate. We present a dynamic structural framework to estimate the returns to seller reputations in freelance sites. In our model, each period, a buyer decides whether to choose a bid from her current set of bids, cancel the auction, or wait for more bids. In the process, she trades-off sellers' price, reputation, other attributes, and the costs of waiting and canceling. Our framework addresses 'dynamic selection', which can lead to underestimation of reputation, through two types of persistent unobserved heterogeneities – in bid arrival-rates and buyers' unobserved preference for bids. We apply our framework to data from a leading freelance firm. We find that buyers are forward-looking, that they place significant weight on seller reputation, and that not controlling for dynamics and selection can bias reputation estimates. Using counterfactual simulations, we infer the dollar value of seller reputations and provide guidelines to managers of freelance firms.

1 Introduction

Online freelance marketplaces are websites that match buyers of services that can be delivered electronically with sellers or freelancers – self-employed individuals/teams, who offer their services on a per-job basis or for a fixed hourly rate. The most popular freelance marketplaces are Elance, Guru, vWorker, ODesk, and Freelancer, and the most popular categories of jobs are web development, programming, writing, translation, design, and multimedia (Kozierok, 2011). These websites typically use auction mechanisms to match buyers and sellers, though their mechanism differs from traditional auctions in three important ways. First, they follow a reverse auction format, where buyers post jobs and sellers bid for jobs. Second, buyers don't wait for the auction to end to make decisions; rather, in every period they decide whether to terminate the auction (by choosing a submitted bid or canceling the auction) or continue waiting. Third, the lowest priced bidder is not the default winner; rather the buyer chooses the winner based on her discretion, and in doing so she may trade-off sellers' reputations, bid prices, other bid attributes, and the cost of waiting and canceling.

Online freelancing has grown tremendously in the last few years. Industry revenues in 2010 were over \$360 million, with a 61% increase from Q1 to Q4 (Morgan, 2011). This surge can be attributed to two factors. First, technological innovations such as electronic deliverability of jobs and fast Internet connections have increased the supply of jobs that can be performed by freelancers. Second, freelance markets offer a low cost way for geographically distant players to trade, especially since there is an abundance of unemployed skilled workers in emerging economies (Indian subcontinent, Eastern Europe) that have low costs of living, and a healthy demand for skilled workers in developed countries, where local labor is expensive.

Despite this recent growth, online freelancing faces many challenges, the primary one being the classic *information asymmetry* or lemon's problem (Akerlof, 1970). Note that buyers face considerable risks in these marketplaces – sellers may deliver low quality services, abscond with advance payments, hold-up the job without completing it and/or delay it, steal Intellectual Property (IP) given to them during the job and sell it to a competitor or use it themselves. While it is theoretically possible to contract on quality, service contracts are notoriously difficult to spell out and enforce (Brousseau and Glachant, 2002). Even if contracts could be made, since most sellers are geographically distant from buyers and belong to developing countries, where IP rights are comparatively lax and legal systems corrupt (Park, 2008), it is exceedingly difficult for injured buyers to obtain legal restitution. These frictions can preclude most transactions.

Online freelance marketplaces seek to mitigate these risks through reputation mechanisms designed to decrease the information asymmetry between players.[1] Typically, reputation systems follow a two-way feedback mechanism – after each transaction, both the seller and buyer are allowed to numerically rate each other, and these ratings are made available to their prospective clients. Intuitively, these feedback systems are designed to incentivize players to behave well in the current period using the threat of future punishment.

However, there is no clear consensus on the effectiveness of these reputation mechanisms. They can fail due to many reasons – imperfect monitoring (Holmstrom, 1999; Cripps et al., 2004), cheap identities

---

[1]Some online freelance marketplaces also offer escrow and arbitration. Escrow solves the problem of what comes first – work or payment. Arbitration can partially solve the problem of quality control if the buyer and seller can pre-agree to abide by the site's decision in case of future differences. However, both escrow and arbitration do not solve the fundamental information asymmetry problem; they simply allow for better distribution of risk among the buyer, seller, and freelance site. So any site using these two mechanisms exclusively would have to charge high commissions to cover its costs, thereby precluding low-value transactions.

(Friedman and Resnick, 2001; Dellarocas, 2003), reliance on unverifiable and voluntary feedback (Resnick and Zeckhauser, 2002), vulnerability to Sybil attacks (Douceur, 2002), and retaliation concerns (Cabral and Hortacsu, 2004), to mention a few. While some solutions have been proposed to these problems (Jøsang et al., 2007), few of which have also been adopted by freelance sites (see §3.1), it is not clear how far they go in building robust reputations. Given that the entire freelancing industry is sustained by feedback-based reputation systems, a good understanding of their effectiveness has implications for the users of the freelance sites, freelance sites themselves, and policy makers interested in regulating this industry. However, to our knowledge, there exists no systematic analysis of returns to seller reputation in freelance marketplaces.

This gap in the literature largely stems from the two key challenges involved in estimating returns to seller reputations in freelance marketplaces. First, in this setting, we need to account for the *option value of waiting* because buyers face dynamic considerations. Figure 1 shows the relationship between the exit period of the buyer and her probability of choosing a bid in our data. Note the unmistakable downward trend – buyers who wait longer are more likely to cancel the auction rather than pick a bid. One reason for this could be that buyers are forward looking, *i.e.*, only those buyers who receive a mediocre set of bids remain in the system (hoping to receive better bids), many of whom eventually cancel as good bids fail to materialize. Hence, ignoring these dynamics can bias the estimates of reputation.

Second, we need to account for *dynamic selection*. The downward trend in Figure 1 could also be due to selection bias – buyers who repeatedly wait are a self-selected group and their persistent tendency to not choose a bid could be due to unobservable (to the researcher) factors that affect their utility from bids. For instance, some buyers may have good unobserved outside options (other freelance sites, own coding abilities), while others may specify their auctions poorly and the bids they attract may not satisfy their requirements. Yet others may have private information on the number of bids they expect to receive in future and/or have an inherently low taste (value) for bids. There may also be significant differences in the average unobserved quality of the bids received by buyers. Such persistent unobservables lead to dynamic selection – the surviving buyers in any period are not random, but are a self-selected group that have low unobserved taste for bids, and thus are more likely to cancel the auction than early exiters. Not controlling for dynamic selection can bias the estimates of reputation; it can lead to systematic underestimation of reputation if a significant fraction of buyers have low unobserved preference for bids.

These challenges cannot be addressed by hedonic regressions and static discrete choice models because they ignore both inter-temporal trade-offs and dynamic selection. Dynamic models without persistent un-observed heterogeneity can handle inter-temporal trade-off, but not selection. In our setting, we find that a static model considerably overpredicts cancellation for early deciders (by 49.92%) and underpredicts cancellation for late deciders (by 44.9%). Similarly, we find that a dynamic model without persistent unobservables underpredicts bid choice in the earlier periods (by 34.71% in period 1) and overpredicts it in the later periods (by 44.06% in the 13^{th} period). Thus, these models not only furnish biased estimates of reputation, but are also very poor in terms of fit.

In this paper, we present a structural framework to estimate the returns to seller reputations in reverse auction settings that can accommodate both dynamics and self-selection. In our partial equilibrium framework, we model buyers' decisions while taking sellers' behavior as given. Buyers face uncertainty over the

number of bids they expect to receive in the future and attributes of those bids (price, seller reputation, *etc.*). Each period, they solve a dynamic programming problem to decide whether to terminate the auction (by choosing one of the submitted bids or canceling) or continue to wait for another period. Our model allows for two types of persistent unobserved heterogeneities – in bid arrival rates across auctions and in buyers' unobserved preference for bids.

Estimation of the model is complicated by the size of the state space and the presence of persistent unobservables. In general, high-dimensional state spaces are intractable with nested fixed point algorithms (Rust, 1987) and are estimated using computationally light two-step methods (Hotz and Miller, 1993). However, our state space is intractable even with standard two-step methods. Moreover, two-step methods have traditionally suffered from their inability to account for persistent unobservables (Aguirregabiria and Mira, 2010). To address these issues, we adapt the two-step estimation framework recently developed by Arcidiacono and Miller (2011) as follows. First, to ensure computational tractability, we reformulate our value function by exploiting the finite-dependence properties of our data. Second, we employ an augmented EM loop that nests a two-step estimator, *i.e.*, we recursively calculate and update the Conditional Choice Probabilities (CCPs) and structural parameters till convergence. A key issue with the use of two-step CCP-based methods in models with persistent unobserved heterogeneity is the non-parametric identification of CCPs. While general proofs that allow unrestricted state transitions are available (Kasahara and Shimotsu, 2009), they are not applicable to our setting. So we derive the conditions for non-parametric identification of CCPs and state transitions in reverse auction settings. Finally, we discuss the identification of the unobserved types and the exclusion restrictions that allow us to empirically estimate the discount factor.

We apply our empirical framework to data from a leading online freelance firm and present six key findings. First, we find that buyers place significant weight on seller reputations. Buyers not only value sellers with high average ratings, but also those with a large number of ratings. The returns to reputation manifests itself in higher probabilities of being chosen as well as the ability to charge higher prices. Second, we find that buyers prefer sellers with low bid prices, sellers with whom they have interacted in the past, and sellers from developed countries. Third, we find that about 30% of the buyers have high unobserved value for the jobs posted, while 70% have low unobserved value. Fourth, we find that not controlling for dynamics and persistent unobserved differences between buyers can lead to serious biases in the estimates of reputation. Fifth, we find that, on average, 87% of buyers considering entry actually choose to enter the market, and that this number varies significantly with their unobserved type. Finally, we estimate the daily discount factor to be 0.88, *i.e.*, buyers are forward looking, but impatient, especially when compared to the discount factor implied by yearly interest rates. Our finding highlights the importance of estimating (as opposed to assuming) the discount factor, especially in banking-unrelated settings.

Next, we present results from a series of counterfactual experiments that quantify the impact of regime changes on auction cancellation rates and site revenues. We do note that because we have a partial equilibrium model, these results are contingent on the assumption that sellers' side behavior remains the same. First, we switch off the site's reputation system and find that site revenues fall by 11.1%; and that revenue loss is the highest from high value auctions. This suggests that the site's reputation system is a significant source of revenue for it. Second, we find that increasing the supply of sellers lowers cancellation rates, but

has no significant impact on revenues! This is because a higher supply of sellers decreases average transaction prices and hence the site's commissions. On the other hand, increasing the fraction of high reputation sellers, while keeping supply constant, has the opposite impact – cancellation rates remain the same, but revenues increase. This is because buyers now replace low reputation, low price sellers with high reputation, high price sellers, thereby driving up the transaction prices and the site's commissions. Together, these two results provide valuable guidelines to managers of freelance sites. First, increasing the supply of sellers uniformly can negatively impact revenues. So decreasing commission rates across the board to either attract sellers from other sites or incentivizing all existing sellers to bid more is not a profitable strategy. Second, it is important to incentivize high reputation sellers alone to bid more and win auctions at higher prices. Incentive mechanisms that selectively lower commission rates for high reputation sellers and/or provide better services to them are recommended.

Finally, we examine whether the site can benefit from charging buyers a fee to post an auction. We consider two types of fees – fixed fees and a percentage of maximum bid. Auction fees have two opposing effects on revenue. On the positive side, the site has a new revenue stream. On the flip side, some buyers who might previously procured from the site now do not even enter the auction. This leads to lower revenues from commissions. These two opposing forces give rise to an inverted U-shaped curve. Specifically, we find that – a) fixed auction fees dominates auction fees based on % of MaxBid and b) site revenues are maximized at an auction fee of approximately $2.75.

In sum, our paper makes three key contributions to the literature. First, from a methodological perspective, we provide a dynamic structural framework to model and estimate the value of bidder attributes in reverse auctions. Our framework not only allows for large state spaces, but also for the option value of waiting and dynamic selection. The framework is fairly general and can be adapted to a large class of optimal stopping problems, *e.g.*, generalized search models, procurement auctions with unspecified end dates, and dating or marriage decisions. Second, from a substantive perspective, we quantify the returns to reputation and other seller attributes in freelance markets. We also estimate the extent of buyer impatience and the distribution of persistent unobservable types in these markets. As far as we know, this is the first paper in Marketing to study freelance markets. Since these markets are becoming increasingly popular, we believe our substantive findings will be of value to researchers, managers, and policy makers interested in this area.

Third, from a normative perspective, our work offers guidelines to sellers and managers of freelance sites and policy-makers. From a seller's perspective, our estimates of buyer utility clarifies how buyers' trade-off reputation and price, and can help them optimize their efforts towards improving own reputation and bidding strategies. From freelance sites' perspective, our framework can be used to gauge the effectiveness of their current reputation systems and evaluate the value of implementing a more robust one. Moreover, our counterfactual results can help sites design better incentive mechanisms for their members. Finally, from policy-makers' perspective, online freelancing is a large and growing industry that contributes significantly to offshore outsourcing of jobs, thereby putting it at the center of the raging debate on the impact of offshore outsourcing on local economies (Mankiw and Swagel, 2006; Lacity and Rottman, 2008). While a complete analysis of the costs and benefits of offshore outsourcing is outside the scope of this paper, our estimates can serve as inputs in the larger cost-benefit analysis that policy-makers must undertake to settle this debate.

2   Related Literature

Our paper relates to four broad streams of literature.

First, our paper relates to the growing literature on online auctions in marketing. Park and Bradlow (2005) and Bradlow and Park (2007) examine bidder behavior in online auctions, while Zeithammer (2006, 2007) examine the dynamics of optimal bidding strategies and buyer behavior with forward-looking sellers. Zeithammer and Adams (2010) investigate the validity of the ubiquitous assumption that online sealed bid auctions are strategically equivalent to second-price auctions. Finally, Yao and Mela (2008) consider a model of both seller and buyer behavior and estimate the impact of varying commission rates and the value of sellers to the marketplace.

Second, it relates to the literature on eBay auctions that uses hedonic regressions to evaluate the value of reputations in eBay (Kalyanam and McIntyre, 2001; Eaton, 2002; Jin and Kato, 2002; Melnik and Alm, 2002; Cabral and Hortacsu, 2004; Bajari and Hortacsu, 2004; Lucking-Reiley et al., 2007). Our paper differs from this research both substantively and methodologically. Substantively, we study seller reputations in a marketplace for services rather than goods, and our setting involves reverse auctions as opposed to traditional auctions. Buyers of services face higher risks, because unlike physical goods, services often have no external brand value, no physical attributes that can be shown in photos, or third-party valuations. Seller reputations may therefore play a more central role in transactions of services. The reverse auction mechanism is not only a substantive difference, but also presents methodological challenges that cannot be addressed by hedonic regressions – buyers face dynamic considerations in this setting and can exit the auction at each period, giving rise to dynamic selection. It is also difficult to interpret the results from hedonic regressions as buyer valuations or some other primitive construct unless we make strong (and unrealistic) assumptions regarding the auction setting (Bajari and Hortacsu, 2004; Rezende, 2008). In contrast, our structural framework based on the utility-maximization framework is capable of handling both dynamics and selection. Moreover, our results can be directly interpreted as primitives that determine buyer utilities and therefore be used to conduct counterfactual experiments to evaluate the impact of policy changes.

Third, our paper relates to the sequential search literature, where an agent who has imperfect information on a set of alternatives sequentially searches for the best option among them by paying fixed cost for each search. Empirical models of sequential search endogenize choice sets in demand estimation and allow researchers to estimate search costs, furnish better estimates of price elasticity, and explain price dispersion for homogeneous goods (Hong and Shum, 2006; Koulayev, 2009). Our model differs from search models in two important respects. First, the sequence of bid arrivals in our model is not determined by the buyer, so unlike search models, sequence provides no information on the reservation prices of bids. Second, in search models, agents have perfect information on the observed attributes of the products and are assumed to be searching only for the realization of the error term (Kim et al., 2010). Together, these two assumptions ensure that search models have an analytical solution (Weitzman, 1979), and hence estimation does not involve numerical solutions to the Bellman equation. This greatly simplifies the modeling challenges. However, in our case, it is unreasonable to assume that buyers have perfect information on the attributes of future bids; in fact buyers don't even know the number of bids they will receive in future periods. We thus have to specify and estimate a full-blown dynamic discrete choice model.

5

Fourth, our paper relates to the small but growing literature that controls for dynamic selection in dynamic discrete choice contexts, *e.g.*, Arcidiacono (2005) in college admissions, Carro and Mira (2006) in couples' contraception and sterilization decisions, and Arcidiacono et al. (2012) in teenage sex and contraception choices). It also relates to the broader literature on finite mixture models. Starting with Dempster et al. (1977), researchers in a variety of fields have employed finite mixtures to accommodate latent unobserved heterogeneity. In the marketing literature, finite mixtures were pioneered by Kamakura and Russell (1989) and Chintagunta et al. (1991), and have since been used extensively. See McLachlan and Peel (2004) and Allenby and Rossi (1998) for detailed discussions of finite mixture models.

3 Setting and Data

3.1 Setting

Our data comes from a leading online freelance firm that had over 320,000 registered freelancers and 150,000 registered buyers, as of March 2011. Membership is free and there is no fee for either posting an auction or for bidding. The site receives about 400 new auctions every day, a vast majority (over 80%) of which are technology-oriented. It follows a sealed-bid reverse auction format, *i.e.*, sellers have no information on other bids received by the buyer. Our data comprises a random sample[2] of all unabandoned[3] public auctions that have two-week expiry periods and maximum bids in the range of $10 to $100 (in increments of $10) initiated from January 1st 2008 to December 31st 2010.[4]

We now describe the auction process in greater detail below.

- A buyer with a procurement need initiates an auction with two key pieces of information:
  - Project title and description – Most project descriptions are very short and generic. There are two reasons for this. First, it is costly (in time and effort) to write out all the project details and most buyers find it easier to talk to the winning bidder by phone or email after the auction. Second, auctions postings are visible to all site members. Therefore revealing project details in the posting carries privacy risks. See Table 11 for examples of project titles and descriptions.
  - MaxBid – maximum amount that the buyer is willing to pay for the project. While providing a MaxBid is optional, most serious buyers choose to do so because it conveys information to the seller about the size and difficulty of the project. While buyers can submit any MaxBid they want, most of them choose multiples of 10. In our analysis, we exclude all auctions where MaxBids are not multiples of 10 because we need to estimate non-parametric joint distributions of bid attributes for each MaxBid included in the analysis in order to model buyers' expectation of future bids' attributes (see §5.4.2). Numbers that are not multiples of 10 have very few auctions that specify them as the MaxBid, making it difficult for us to generate bid distributions for them.

---

[2]We do not reveal the size of the full dataset in order to preserve the privacy of the firm.

[3]Some buyers start auctions, but don't monitor them or take any actions subsequently. We exclude such auctions from our analysis because we do not know whether the buyer was actually making any decisions or not, and if so, when exactly the buyer chooses to ignore the auction. Estimates from a model including these auctions are qualitatively similar to those presented here.

[4]The site also allows auctions where sellers bid for the hourly rate to be paid to them while they work on the project. This is in contrast to the pay-per-project auctions that we study, where sellers bid on the payment for completing the project. Usually, projects that can be well-defined ahead of time are sold on a pay-per-project basis, while projects with undefined scope are sold on a pay-per-hour basis. However, the pay-per-hour format was only recently introduced by the site and forms a very small portion of their business. So we exclude them from our analysis.

- After confirming that the project does not involve illegal activities, the site posts the auction on its public forum, which can be browsed by all its members. Sellers can also obtain up-to-date information on new project postings by subscribing to newsletters from the site. The posting contains information provided by the buyer (project description, auction start and expiry date, and MaxBid) as well as information on the buyer herself (*e.g.*, her past ratings and her geographic location) through a link to her homepage.
- Sellers can start bidding for the job as soon as the auction goes live. The bid along with the seller's past average rating becomes available to buyers through the site. The bid also contains a link to the seller's homepage which has additional information on her.
- The buyer can stop the auction at any point in time by either picking one of the bids she has already received (if any) or canceling the auction. If she hasn't picked a bid after two weeks, the auction is automatically canceled.[5]

The site employs a mandatory escrow and offers free arbitration services for these auctions. It charges a fees for its services, in the form of a percentage commission on the transaction amount, which is paid by the winning bidder. For example, if a seller with a bid of $50 wins a project, the buyer escrows $50 with the freelance site, and after the project is completed, the site releases $50 minus its commission to the bidder.

We now describe the reputation system used by the marketplace in detail. The site uses a symmetric numeric rating scale of 1-10 (and optional text comments) for both buyers and sellers. A rating of 1 stands for very bad and 10 for excellent. The site has implemented the following measures to make the reputation system robust. First, after a trade, both the buyers and seller are given a fixed time period to rate each other, after which they lose the right to rate. The ratings are revealed publicly only after both parties have rated each other. If one of the parties fails to turn in its feedback, then the other party's rating is revealed only after the fixed time period, at which point the delinquent party cannot retaliate. Second, if there is a dispute following the trade and the case goes into arbitration, both parties lose the right to rate each other, though the neutral arbiter may rate either or both players. Third, members' homepages have information on all their past trades, *i.e.* members cannot selectively hide ratings. Overall, the site has a relatively high feedback rate compared to sites such as eBay, where the feedback rate is barely 50% (Resnick and Zeckhauser, 2002). In our data, 92.4% of the buyers have rated their sellers, while 71.8% of the sellers have rated their buyers.

3.2 Data

For each auction in our sample, we have the following information:

- The MaxBid of the auction.
- Start and end dates of the auction.
- Number of bids received for each day the auction is active.
- The following buyer attributes:
  - Geographic region of the buyer – region codes are shown in Table 1.
  - Total number of past auctions initiated by the buyer.
  - Cancel ratio – fraction of past auctions that the buyer canceled. A buyer who has initiated 10

---

[5]A winning bidder has 24 hours to reject the job without penalty. In such cases, the buyer may pick another bidder. In the data, we observe very few occurrences where winning sellers reject a won job. In those cases, we treat the first choice of the buyer as the winning bid. The qualitative results remain unchanged if we instead drop these auctions from our analysis.

auctions and picked winners in 7 auctions, has a cancel ratio of 0.3; by default, it is zero for buyers with zero past auctions. Cancel ratio is indicative of the buyer's inherent choosiness and/or the quality of her outside options.

- Number of past ratings and the sum of all past ratings.
- Mean rating – defined as the 'sum of all past ratings/total number of past ratings' if the buyer has at least one rating, and zero otherwise.
- Tenure on the site – number of days since the buyer signed up.

- The following attributes for all bids received – bid price, seller's geographic region (Table 2), number of her past ratings, sum of her past ratings, her mean rating, and an indicator for whether she has worked for the buyer in the past on this site.[6]

Table 2 provides an overview of the number of auctions in our data, by MaxBid, and their outcomes. Note that only 78.23% of the auctions in the data end with the buyer picking a bid, while the rest are canceled. On average, auctions are active for about 2.67 periods and the average lifespan of an auction increases with MaxBid (Table 5). That is, buyers who value the job more or those with larger projects seem to wait longer. Nevertheless, even within the same MaxBid, there is significant variation in when buyers close their auctions. An average auction receives about 9.54 bids, with the median being 6 (Table 5). Because this number does not take into account the number of periods the auction was active, we present Figure 2, which shows the distribution of the average number of bids received per period by the auctions in our data. There is considerable heterogeneity across auctions in the average number of bids received per period and the distribution in Figure 2 exhibits a long tail – a majority of the auctions (50.76%) receive no more than an average of two bids per-period, while a few of them (18.06%) receive ten or more.

Table 3 shows the summary statistics of buyer attributes.[7] The median buyer has about 10 ratings, with an average rating of 9.96. However, a big chunk of them (15.1%) have no past ratings, and while another significant chunk (17.49%) have a mean rating of 10, with 10 or more ratings. Most buyers in the data have previous experience on the site; the median buyer has posted 11 successful auctions (in which she picked a bid) and 7 canceled auctions. Finally, as shown in Table 4, the majority (81.88%) of buyers belong to developed English-speaking countries, *i.e.*, Region 2.

Table 6 shows the summary statistics of seller attributes. A large percentage (24.87%) of bidders have no ratings, and 20.97% of them have a very good reputation (with a mean rating of 10 with 10 or more ratings). About 1.98% of the bidders have interacted with the buyer in the past. Unlike buyers, majority of the bidders (56.69%) belong to the Indian sub-continent (Table 4). The other three regions are about equally represented (approx. 14-15% each). Moreover, the distributions of the sellers' geographic region varies with the buyer's region. Finally, while a large number of bidders quote the MaxBid as their price (40.34%), many of them also quote much lower prices, with 4 being the lowest observed quote (Table 7).

Table 6 also shows the summary statistics of accepted bids. There are systematic differences in the attributes of winning bidders compared to the full distribution of bidders. On average, they quote lower

---

[6]Both buyer and seller reputation metrics may change during the course of the auction if they receive new ratings. However, such changes are very small and/or infrequent. So for each auction, we obtain all buyer attributes at the beginning of the auction, and for each bid, we obtain seller's attributes at the date of bid submission, and then treat them as constant for the duration of the auction.
[7]In order to preserve the privacy of buyers and the freelance site, summary statistics for buyer tenure are not shown.

prices (see Table 7), have significantly better reputations, are more likely to belong to developed countries, and have a higher likelihood (10.46%) of past interaction with the buyer.

4   Model-free Evidence

We now present some model-free evidence in support of the effectiveness of the reputation system. First, note that, only 40.96% of the buyers who eventually pick a bid (*i.e.*, don't cancel the auction), choose the lowest priced bid (Table 2). This tendency is also more pronounced for larger projects (higher MaxBids). For example, while 52.48% of the buyers in the category MaxBid = 10 pick the lowest priced bid, only 37.83% in the category MaxBid = 100 do the same. These patterns in the data suggest that considerations other than price, potentially sellers' reputations, play an important role in buyers' decisions.

Second, we examine whether there is sufficient variation in the distributions of equilibrium prices. When reputation systems fail, high quality or truthful sellers desert the market, leaving only low quality sellers, who in turn receive low prices in equilibrium (Akerlof, 1970). In contrast, a healthy marketplace with an informative reputation system can support both high and low quality sellers, with the former receiving a premium for reputation. So the absence of price variation in a market usually indicates the failure of its reputation system. Table 7 gives the summary statistics of the trading price for all the MaxBids and Figure 3 shows the distribution of the prices received by sellers in equilibrium for MaxBid = 100. Together, they confirm that equilibrium prices are reasonably diffuse.

Third, we check whether buyers are more willing to pick a high reputation seller compared to a low reputation seller. Figure 4 presents the CDFs of the number of past ratings for two sets of sellers – (a) sellers whose bids were chosen by the buyer, and (b) sellers whose bids were not chosen. Note that there is a clear difference in the CDFs – on average, winning bidders have more ratings than the losing bidders. Figure 5 presents the distributions of the mean ratings for the same two sets of sellers. Here, we find that, on average, chosen bidders have a better mean rating than those who weren't. Taken together, these findings suggest that buyers value and prefer sellers with good reputations.

Finally, we examine whether buyers are willing to pay higher prices to high reputation sellers, *i.e.*, are reputations and equilibrium prices positively correlated? In Figure 6, we present the distributions of prices, for chosen bids, for three types of sellers – (a) sellers who have no past ratings, (b) sellers who have some reputation, and (c) sellers who have an excellent reputation. On average, sellers with no past ratings receive the lowest prices, while those with medium reputations receive somewhat higher prices, and those with the highest reputations receive the best prices.

While these preliminary findings are suggestive rather than conclusive, they nevertheless support the existence of positive returns to reputation for sellers in this market.

5   Model

We now specify and estimate a formal model of buyer behavior that rationalizes the choices observed in the data. We now briefly summarize our empirical framework. First, in §5.1, we present the model set-up and timeline. Second, in §5.2, we specify the set of variables that affect buyers decisions. We also specify which of these vary with time, which remain constant, and which of these are observable to researcher. Next, in §5.3, we present buyers' flow utilities or per-period utilities as functions of observable and unobservable

9

state variables. Then, we dive into the dynamics of the problem. In §5.4, we explain how we capture buyers' future expectations. We model buyers' beliefs on the number of bids they expect to receive using Fixed Effects Poisson in §5.4.1. To model buyers' beliefs on the attributes of the bids they expect to receive in future, we use a combination of nonparametric joint distributions (see §5.4.2) and multinomial logit models (see §5.4.3 and §5.4.4), depending on whether the bid attribute is continuous or discrete. Finally, in §5.5, we combine all these pieces to formulate buyers' dynamic optimization problem.

5.1 Set-up

Time is discrete and indexed by $t = \{1, 2, ..., T\}$. There are N buyers, who are indexed by i and have defined preferences over a sequence of states of the world from $t = 1$ to T. Each buyer initiates one auction and with some abuse of notation, we often use buyer i and auction i synonymously. Auctions go live at the beginning of period 1 and buyer i makes her decision d_{it} at the end of every period, till she makes a terminal decision or till the auction expires at T. In our data and analysis, one period is equivalent to a day and the expiry period is $T = 14$.[8] Each period, the buyer's options are – cancel the auction, pick one of the bids she has already received (if any), or wait for another period. The former two decisions are terminal, *i.e.*, once she cancels or picks a bid, she makes no more decisions. The option of waiting is unavailable at T, whence the buyer has to either choose one of her bids or cancel. i's alternatives are indexed by j. For periods $1 \leq t \leq T - 1$, $j \in \mathcal{J}_{it} = \{1, 2, ..., J_{it}\}$, where alternative $j = 1$ refers to wait, $j = 2$ to cancel, and $j = \{3, ..., J_{it}\}$ to the bids that have arrived so far. For period T, the alternatives are $\mathcal{J}_{iT} = \{2, ..., J_{iT}\}$. Finally, note that the buyer faces a non-stationary dynamic optimization problem because the horizon is finite; there is a clear terminal period T.

5.2 State Variables

Let $\{x_{it}, \zeta_{it}\}$ denote the state of the world for buyer i at period t, where x_{it} is the set of observable and ζ_{it} is the set of unobservable (to the researcher) state variables.

5.2.1 Observed State Variables

We consider two sets of observable state variables that influence the buyer's decision – time invariant and time varying. The time invariant state variables are – (a) m_i, MaxBid specified by the buyer, (b) n_i, number of ratings received by the buyer, (c) r_i, mean rating of the buyer, (d) a_{si}, number of past auctions in which i picked a bidder (*i.e.*, didn't cancel), (e) a_{ci}, the number of past auctions canceled by i, (f) g_i, buyer i's geographic region, and (g) l_i, length of buyer tenure on site. These time invariant state variables remain constant through the auction, *i.e.*, they transition into the same values every period.

The time varying state variables are time t, the total number of bids received so far, B_{it}, and a $B_{it} \times 5$ matrix of bid attributes. The j^{th} row of this matrix captures five key attributes of the j^{th} bid received by i – (a) bp_{ij}, bid price, (b) bn_{ij}, number of ratings received by bidder j, (c) br_{ij}, mean rating of j, (d) bg_{ij}, bidder j's geographic region, and (e) bt_{ij}, is an indicator variable that is 1 if i and j have traded in the past.[9]

---

[8]In practice, time is continuous and a buyer may make decisions at any point in time, as many times as she wishes, till she makes a terminal decision or till the auction ends. However, we don't observe when buyers make decisions. While we have data on when a terminal action was taken, we have no information on the intermediate decision points at which the buyer chose to wait. We therefore aggregate the data over 24-hour intervals and assume that all buyers make their decisions on a daily basis.

[9]For past interactions, it is possible to include the rating given to j by i as another state variable too. However, in general, we

5.2.2 Unobserved State Variables

Time invariant unobserved state variables – These variables capture the persistent unobserved hetero-geneity between buyers. We consider two types of persistent unobserved state variables – $\{\eta_i, s_i\}$. η_i is a time-invariant buyer (or auction) fixed effect that impacts the number of bids received by buyer i, though it does not affect i's flow-utility.

$s_i = \{1, ..., S\}$ denotes i's unobserved type, drawn from a set of finite types. s_i captures the unobserved aspects of the buyer, auction, and bidders that influence a buyer's tendency to pick a bid. For example, some buyers may have better outside options through other freelancing sites or own coding abilities. Such buyers would exhibit a persistent tendency to cancel or wait, and not choose bids. Or, a buyer may specify her auction poorly and the bids she attracts may not satisfy her requirements, in which case she would persistently avoid choosing a bid. There might be significant differences in the unobserved quality of the bids across auctions, *e.g.*, some buyers/auctions may receive bids that, on average, have lower unobserved quality, and therefore exhibit a persistent tendency to not pick bids. s_i captures all these factors because it links buyers' choices over time and thereby addresses the *dynamic selection* problem discussed in §1. In our application, we allow for two unobserved types: Low type, which has low unobserved preference for bids ($s_i = 1$) and High type, which has high unobserved preference for bids ($s_i = 2$).

With sufficiently long panels for buyers, it is possible to avoid persistent unobservables and instead estimate buyer-specific CCPs, à la Misra and Nair (2011). However, we don't have sufficiently long panels for a large fraction of buyers to adopt this approach. Moreover, this method would not allow us to control for auction-specific persistent unobservables.

Time varying unobserved state variables – ϵ_{it} is a $J_{it} \times 1$ vector with support $R^{J_{it}}$, whose j^{th} component forms the mean-zero additive error to buyers' per-period utility from alternative j. We assume that the errors ϵ_{it} are i.i.d over time and drawn from a Generalized Extreme Value (GEV) distribution, which yields nested logit probabilities in a static setting. All the bid options are in one nest, and the cancel and wait options are in two separate singleton nests. Let $\sigma \in [0, 1]$ be the correlation of errors in the nest with the bid options, where $\sigma = 0$ implies perfect correlation and $\sigma = 1$ indicates no correlation. Errors across nests are not correlated. Assuming a GEV distribution considerably eases the computational burden because of two reasons. First, it gives us closed-form expressions for choice probabilities (McFadden, 1978; Rust, 1987). Second, it allows us to derive an analytical relationship between choice probabilities and the future value function (Hotz and Miller, 1993; Arcidiacono and Miller, 2011).

5.3 Flow Utility

There are two possible ways to model buyers' flow utility – either derive it from a micro-model of buyer and seller behavior based on assumptions on economic primitives (*e.g.*, distribution of seller types and buyer's rating policies and learning behavior), or choose a convenient parametrization that is flexible enough to capture the patterns in data. While the former is more theoretically appealing, it has three main drawbacks (Berry and Reiss, 2007; Ellickson and Misra, 2011). First, it requires much more structure and additional

---

find that sellers don't bid on auctions by buyers who gave them low ratings previously. So the ratings from past interaction don't provide much more information than the indicator of past interaction.

data on the primitives of the market and information on repeated interactions between buyers and sellers, which are not available to us. Second, it is likely to become unwieldy in most real-world settings such as ours, rendering estimation intractable. Third, without information on true seller types, bidding strategies, and buyer's rating behavior, identification of primitives will still be driven by functional form assumptions. Hence, we adopt the latter approach in our paper.

Let buyer i's present discounted value of her lifetime utilities at period t from making decision $d_{it} \in \mathcal{J}_{it}$ be $\Sigma_{k=t}^{T} \delta^{k-t} U(d_{ik}, x_{ik}, s_i, \epsilon_{ik})$, where $\delta \in (0, 1)$ is the discount rate and $U(d_{it}, x_{it}, s_i, \epsilon_{it})$ is buyer i's flow utility in period t. $U(d_{it}, x_{it}, s_i, \epsilon_{it})$ is additively separable as follows:

$$U(d_{it}, x_{it}, s_i, \epsilon_{it}) = u(d_{it}, x_{it}, s_i) + \epsilon_{it}(d_{it}) \tag{1}$$

where $\epsilon_{it}(d_{it})$ is the d_{it}th component of the $J_{it} \times 1$ i.i.d GEV error vector ϵ_{it}.

There are two points of note here. First, flow utilities are not functions of the unobserved time invariant state variable η_i, because buyers don't receive any instantaneous utility in anticipation of future bids. Second, time does not enter the flow utility as a state variable, *i.e.*, we do not assume duration dependence. There are two reasons for this – a) duration dependence in consumption usually comes into play when consumers accumulate stocks of utilities over time through experience goods such as vacations or recreational golf (Hartmann, 2006). In our context, there is a one time exit decision, and there is no rational underpinning for a utility stock from waiting. b) More importantly, we cannot separately identify persistent unobserved heterogeneity in buyers' preference for bids (s_i) and duration dependence in utilities because we don't have data on multiple auctions for a large fraction of buyers. Of course, optimal decisions will be functions of η_i and time through the dynamics of the problem.

We now specify the deterministic components of buyers' per-period utility from waiting ($j = 1$), canceling ($j = 2$), and picking a bid ($j \geq 3$). In all discrete choice models, only differences in utilities matter. We therefore need to normalize the utility from one alternative to zero and the utilities from all other alternatives are specified in relation to this option. While we can choose any alternative as this base choice, we choose cancel as that option for a key reason – it gives us finite-dependence, *i.e.*, it gives us a decision that resets the state-space to a known configuration and has a well-defined utility. As Arcidiacono and Ellickson (2011) point out, a clever choice of a base decision that exploits finite dependence can greatly simplify estimation. See §6.1 for details on how this normalization simplifies estimation.

$$u(2, x_{it}, s_i) = 0 \tag{2}$$

Next, we specify the flow utility from waiting (in comparison to that from canceling). Canceling allows buyers to close the auction, and finish the work through other means immediately, *e.g.*, do it themselves, hire a local programmer, visit another freelance site. While a buyer who chooses to wait may also explore outside options without any monetary costs (as the site does not charge any fees for keeping an auction active), there are other hassles associated with waiting/maintaining an active auction – every time the buyer signs into the site (even for activities unrelated to this particular auction), she is urged to review bids and close the auction. The site also sends emails and messages with information on new bids, urging buyers to

pick a bid, which some buyers may want to avoid. In our data, we find that most buyers cancel the auction well before the two-week expiry period, which suggests that there are indeed costs to waiting. Therefore, we do not restrict the flow utilities from waiting and canceling to be the same. Rather, we allow the utility from waiting (in comparison to canceling) to vary with buyer and auction specific variables, and let the estimates from the data inform us of buyers' relative preference for both.

$$
\begin{aligned}
u(1, x_{it}, s_i) = W_{iw}(x_{it}, s_i) = W_{iw}(x_{it}) \quad = \quad & \alpha_{w1} + \alpha_{w2}m_i + \alpha_{w3}\ln(n_i \cdot r_i + 1) + \alpha_{w4}I(a_{si} + a_{ci} = 0) \\
& + \alpha_{w5}a_{si} + \alpha_{w6}c_i + \alpha_{w7}l_i + \alpha_{w8}I(g_i = 1) \\
& + \alpha_{w9}I(g_i = 2) + \alpha_{w10}I(g_i = 3)
\end{aligned}
\tag{3}
$$

The flow utility from waiting is allowed to depend on MaxBid (m_i), buyer reputation ($\ln(n_i \cdot r_i + 1)$, ln sum of buyer ratings), buyer's previous experience on the site ($I(a_{si} + a_{ci} = 0)$, an indicator for a new buyer; a_{si}, number of uncanceled past auctions; c_i cancel ratio; l_i length of buyer tenure on site), and her geographic location ($I(g_i = 1)$, $I(g_i = 2)$, $I(g_i = 3)$). W_{iw} is independent of s_i because s_i is defined as i's unobserved taste for picking a bid. By modeling the wait utility as a function of buyer and project specific variables, we are allowing for buyers' cost of waiting to vary with these factors. For example, buyers with larger projects (higher MaxBid) may have lower cost of waiting, while buyers from developing countries may have higher cost of waiting, as they can easily obtain cheap local labor.

Next, we assume that i's utility from picking a bid j (in comparison to canceling) is:

$$
u(\text{bid}_j, x_{it}, s_i) = W_{ib}(x_{it}, s_i) + Y_{ij}(x_{it}) \text{ if } j \geq 3
\tag{4}
$$

$W_{ib}(x_{it}, s_i)$ refers to the aspects of the utility function that are constant across all bids, *i.e.*, they are the same within the nest of bid options. $W_{ib}(x_{it}, s_i)$ is specified as:

$$
\begin{aligned}
W_{ib}(x_{it}, s_i) = \quad & \alpha_{b1} + \alpha_{b2}m_i + \alpha_{b3}\ln(n_i \cdot r_i + 1) + \alpha_{b4}I(a_{si} + a_{ci} = 0) + \alpha_{b5}a_{si} + \alpha_{b6}c_i \\
& + \alpha_{b7}l_i + \alpha_{b8}I(g_i = 1) + \alpha_{b9}I(g_i = 2) + \alpha_{b10}I(g_i = 3) + \alpha_{b11}I(s_i = 2)
\end{aligned}
\tag{5}
$$

It is analogous to W_{iw}, with the unobserved type s_i thrown in. Since we allow for two unobserved types, α_{b11} is the relative preference of High type ($s_i = 2$) for the bid nest, in comparison to the Low type. While s_i enters the flow utility only as an intercept, it should be noted that the tendency to pick a bid is not merely shifted by the intercept, because s_i also enters the utility from waiting (through the value function) in a highly non-linear way. Thus the overall impact of the unobserved type is more involved than it appears at the first glance.[10]

The second part of bid j's utility, $Y_{ij}(x_{it})$, consists of terms that vary across alternatives within the bid

---

[10]In general, because of the inherent non-linearity of dynamic discrete choice models, unobserved heterogeneity is usually included as an intercept; and what may seem like a restrictive formulation in a static model is far from it in a dynamic model. The reasoning for this is very similar to the reason why static logit models suffer from IIA, but dynamic logit models with the same formulation are free from IIA. Please see Carro and Mira (2006); Arcidiacono (2005); Arcidiacono et al. (2007, 2012) for previous examples of this kind of formulation.

nest. Since s_i does not vary within the bid nest, $Y_{ij}(x_{it})$ does not depend on it.

$$\begin{aligned}
Y_{ij}(x_{it}) = {} & \beta_1 \ln(bp_{ij} + 1) + \beta_2 I(bn_{ij} = 0) + \beta_3 \ln(bn_{ij} + 1) + \beta_4 (br_{ij} - \bar{br}) + \beta_5 (br_{ij} - \bar{br})^2 \\
& + \beta_6 (br_{ij} - \bar{br}) \ln(bn_{ij} + 1) + \beta_7 m_i \ln(bn_{ij} + 1) + \beta_8 r_i (br_{ij} - \bar{br}) + \beta_9 b_{ti} \\
& + \beta_{10} I(bg_{ij} = 1) + \beta_{11} I(bg_{ij} = 2) + \beta_{12} I(bg_{ij} = 3) + \beta_{13} I(g_i = 2, bg_{ij} = 1) \\
& + \beta_{14} I(g_i = bg_{ij} = 2) + \beta_{15} I(g_i = 2, bg_{ij} = 3) + \beta_{16} I_3(g_i = bg_{ij} \neq 2) \quad (6)
\end{aligned}$$

$Y_{ij}(x_{it})$ is a concave function of bid j's price, $\ln(bp_{ij} + 1)$, to allow for non-linearities in price sensitivity and the following reputational attributes of the seller – (a) $I(bn_{ij} = 0)$, indicator for no past ratings, (b) $\ln(bn_{ij} + 1)$, a concave function of her total past ratings, (c) $(br_{ij} - \bar{br})$, her mean centered average past rating, (d) its square $(br_{ij} - \bar{br})^2$, and (e) the interaction term $(br_{ij} - \bar{br}) \ln(bn_{ij} + 1)$. (Mean centering aids in the interpretation of parameter values after estimation.) We also allow for interaction between seller's reputation and project size through $m_i \ln(bn_{ij} + 1)$, because buyers' may value seller reputations more for larger projects. Similarly, we allow for interaction between seller's and buyer's reputations through $r_i(br_{ij} - \bar{br})$, because high reputation buyers may place a higher value on seller's reputations. Utility from a bid is also allowed to depend on the past history of the buyer and the bidder because buyers may prefer sellers with whom they have interacted in the past. Finally, we also include geographic region dummies and interaction effects between buyer and seller regions. While the utility function can be tweaked further, we found the above specification to be superior to others that we experimented with.

Note that Y_{ij} only depends on the attributes of seller j, but not her identity. Our specification is similar in spirit to Arcidiacono (2005), who allows students' utilities from attending a specific college to depend on the attributes of the college, but not its identity. We cannot allow for seller identifiers, either through seller dummies or through unobserved seller types, because – a) we don't see enough repeat bids from sellers to identify seller dummies or seller specific unobservable types. b) Attaching an unobserved type to each seller is infeasible given the number of bids in the data. For example, even if we were to allow two unobserved types of sellers, for a typical auction with 10 bids, the number of possible unobserved states would equal 2^{10}. We would then have to integrate out 2^{10} possible states in our likelihood, which is infeasible.

Given that we have included all the key seller specific variables that a buyer observes when making her decision, and given the nature of our setting, where most sellers are small players who have no brand equity beyond the site, this assumption seems reasonable. Moreover, as discussed earlier, even though we cannot control for the identity or unobserved quality of each bid, we are able to control for the average unobserved quality of a buyer's bids through s_i.

5.4 State Transitions

Buyers' decisions not only influence future state transitions, but are also influenced by their beliefs on the evolution of future states. We thus need to model state transitions and specify buyers' beliefs over them. We assume that buyers have rational expectations, *i.e.*, their expectations are consistent with the true state transition probabilities inferred from data.

First, note that, buyer and auction specific state variables (observed and unobserved) remain constant for the duration of the auction. Next, consider the time varying observed state variables. Time transitions

deterministically, increasing by 1 every period. However, bids arrive stochastically and buyers face uncertainty over the attributes of future bids. Bid arrivals are modeled using a fixed effects Poisson. Buyers' expectation on the attributes of future bids is modeled as follows – (a) Price, mean rating, and the number of ratings are modeled using three dimensional non-parametric joint distributions. (b) Indicators for the sellers' geographic region and the past interaction are modeled using two separate logit distributions. Finally, as discussed earlier, the unobserved state variables, ϵ_{it}s, are assumed to be i.i.d over time.

We discuss these state-transition models in detail below.

5.4.1 Poisson Bid Arrival Process

We use a Poisson process to model bid arrival because – a) bids arrive independently in sealed-bid auctions and b) we don't see evidence for over-dispersion in data (Table 5). The Poisson model however needs to capture two important patterns in the data. First, it should allow for auction level heterogeneity since some auctions receive many more bids than others (Figure 2). Further, since only some variation in the number of bids received can be explained by observed buyer and auction characteristics, we also need to account for unobserved heterogeneity in buyers or auctions. Second, the model should capture the fact that bid arrival rates vary over different time periods. For example, we observe that bid arrivals generally slow down with time. To accommodate these considerations, we employ a fixed effects Poisson model.

The conditional probability function of b_{it}, the number of bids received by buyer i in period t, is:

$$h_p(b_{it}|z_{it}, \eta_i, \theta_p) = \frac{\exp(-\eta_i\lambda_{it})(\eta_i\lambda_{it})^{b_{it}}}{b_{it}!} \tag{7}$$

where $\lambda_{it} = \exp(z'_{it}\theta_p)$, η_i is an unobserved buyer-auction specific fixed effect, z_{it} is a set of time varying buyer attributes, and θ_p is the parameter vector to be estimated. Since bids arrive independently and are not visible to other bidders, the number and characteristics of previous bids cannot affect the number of bids arriving in the current time period. Therefore, the only time varying buyer/auction attribute is time t. Hence, we omit the subscript i in z_{it} (and λ_{it}), and define z_t as the set of $T-1$ time dummies. In the Poisson model, all time invariant buyer and auction attributes (including buyer's unobserved type, s_i) are subsumed by the fixed effect η_i. Overall, $\{\theta_p, \eta_1, ..., \eta_N\}$ is the set of parameters to be estimated in this context.

5.4.2 Non-Parametric Joint Distributions of Price, Number of Ratings, and Mean Rating

Our preliminary analysis revealed that the distributions of these three attributes do not follow any specific parametric forms. Therefore, we model their joint distributions using multivariate kernel density functions.

First, we classify all the bids according to the MaxBid (10 to 100) of the corresponding auction. Table 2 shows the total number of auctions and bids in each MaxBid over the observation period. We consider each MaxBid as a separate class because the range and distribution of bid prices is very different across MaxBids (Table 7). Moreover, the correlations between the three bid attributes vary by project size or MaxBid. Next, we sub-classify the bids based on the buyer's reputation in order to capture the differences in buyers' expectations about the bids they might receive in future, as follows: a) sub-class 1: No. of ratings = 0; b) sub-class 2: No. of Ratings > 0 & Avg. Rating < 9.5; c) sub-class 3: No. of Ratings < 15 & Avg. Rating $\geq$ 9.5; d) Sub-class 4: No. of Ratings $\geq$ 15 & Avg. Rating $\geq$ 9.5. Our classification allows buyers with a

large number of high ratings (sub-class 4) to receive better bids (low bid price, high bidder reputation *etc.*),
on average, compared to those with no ratings (sub-class 1). In fact, Koglomorov-Smirnov tests comparing
the distributions for the four sub-classes of bids (for a given MaxBid) reject the null hypothesis of equality of
the distributions. Ideally, of course, we would like to have a more fine-grained sub-classification. However,
further classification is not feasible given the size of our data.[11]

Let $\mathcal{C} = \{1, ..., C\}$ be the set of categories, where $C = 40$ since we have 10 MaxBids, each with four
sub-classes. $c = 1$ denotes MaxBid = 10 and sub-class 1, $c = 2$ denotes MaxBid = 10 and sub-class 2, and
so on. The observed bids in each category are indexed by $q \in \{1, 2, .., M_c\}$, where M_c is the total number
of bids in category c. The q^{th} bid in any category is denoted by the vector $A_q = (bp_q, bn_q, br_q)'$, where the
three elements denote bid price, number of ratings, and mean rating of the seller. We model the probability
density function at a point A in the three dimensional space, in category c, using the multivariate kernel
density estimator:

$$\psi_c(A, h_c) = \frac{1}{M_c \cdot h_c^3 \cdot r(k_c, A)^3} \sum_{q=1}^{M_c} K \left\{ \frac{A - A_k}{h_c \cdot r(k_c, A)} \right\} \tag{8}$$

where h_c is the optimal bandwidth window for category c, $K(\cdot)$ is the three dimensional kernel function
satisfying the property $\int_{R^3} K(A)d(A) = 1$. (We choose the standard trivariate normal in our estimation.)
$r(k_c, A)$ is a scaling parameter that represents the Euclidean distance from A to the k^{th} nearest point in
the data. In the absence of $r(k_c, A)$, the same bandwidth is used for all parts of the distribution. This is
problematic in finite samples because it is difficult to pick one optimal bandwidth for the entire range of the
distribution; low bandwidths lead to spurious noise in the tails of the distribution, while high bandwidths
cause over-smoothing in the main parts of the distribution (Silverman, 1986). Scaling the bandwidth lo-
cally using $r(k_c, A)$ provides a simple but effective solution to this problem. Further, as is common in the
literature, we set $k_c = \sqrt{M_c}$.

The choice of the bandwidth is crucial to the quality of the kernel density estimator. In §6.3.2, we discuss
the estimation of the optimal bandwidth h_c.

5.4.3 Multinomial Logit Model of Bidder's Geographic Region

Sellers can belong to one of four discrete geographic regions (Table 1). Conditional on buyer-specific state
variables and a given draw of bid price, bidder mean rating, and number of bidder ratings, we can model
the distribution of seller's geographic region using a multinomial logit model. Let $h_g(bg_{ij}|gx_{ij}, \theta_g)$ be the
conditional probability of bg_{ij}, where gx_{ij} is the set of state variables that influences the draw of the bidders'
geographic region and $\theta_g = \{\theta_{g1}, \theta_{g2}, \theta_{g3}\}$ are the parameter vectors associated with the regions 1, 2, and
3, respectively. Then, the probability that bidder j in auction i belongs to geographic region q is:

$$h_g(bg_{ij} = q|gx_{ij}, \theta_g) = \frac{e^{gx_{ij}\theta_{gq}}}{1 + e^{gx_{ij}\theta_{g1}} + e^{gx_{ij}\theta_{g2}} + e^{gx_{ij}\theta_{g3}}} \; \forall \; q \; \in \{1, 2, 3\}$$

$$h_g(bg_{ij} = 4|gx_{ij}, \theta_g) = \frac{1}{1 + e^{gx_{ij}\theta_{g1}} + e^{gx_{ij}\theta_{g2}} + e^{gx_{ij}\theta_{g3}}} \tag{9}$$

---

[11]Results are robust to modifications in cut-offs used to sub-classify the data, such as using $n_i > 20$ as the cut-off point between
sub-classes 3 and 4. Moreover, preliminary analysis of the data did not reveal any systematic differences between the distributions
of early and later bids. So we don't further classify bids based on their time of arrival.

Note that $\theta_{g4} = 0$ because $q = 4$ is the base region. In our estimation, we include buyer-specific variables and the three bid attributes drawn from the nonparametric joint distribution in gx_{ij}. The set of parameters to be estimated in this context is $\theta_g = \{\theta_{g1}, \theta_{g2}, \theta_{g3}\}$.

5.4.4 Logit Model of Buyer-Bidder Past Interaction Indicator

The buyer-seller past interaction is characterized using the indicator variable bt_{ij} that is 1 only if the buyer and seller have interacted in the past. Let $h_b(bt_{ij}|tx_{ij}, \theta_t)$ be the conditional probability of bt_{ij} given state variables tx_{ij} and parameter vector θ_t. tx_{ij} consists of buyer-specific state variables and bid price, bidder mean rating, number of bidder ratings, and bidder country. Then, the logit probability of $bt_{ij} = 1$ is:

$$h_b(bt_{ij} = 1|tx_{ij}, \theta_t) = \frac{e^{tx_{ij}\theta_t}}{1 + e^{tx_{ij}\theta_t}} \tag{10}$$

$bt_{ij} = 0$ is the base option, with probability $\frac{1}{1+e^{tx_{ij}\theta_t}}$.

5.4.5 State Transitions and Unobserved Buyer Type s_i

Note that bid arrivals are allowed to be correlated to a buyer's unobserved type (s_i) because the Poisson fixed effect (η_i) subsumes all time invariant buyer attributes, including s_i. We thus allow the number of bids received by a buyer to be correlated to her unobserved preference for bids. For example, a buyer who puts up an ill-specified auction may not receive many bids and also exhibit a tendency to not pick bids.

However, we do not allow the attributes of bids received by a buyer to be correlated to her unobserved type s_i. Recall that we model three of the bid attributes using non-parametric multivariate joint distributions. Allowing for correlations between these bid attributes and unobserved type s_i would require us to specify and estimate non-parametric mixture models at each step of the EM-algorithm, which would non-trivially increase the computational tractability of the model.

5.5 The Buyer's Problem

In each period, buyer i picks a decision that maximizes the present discounted value of her lifetime utilities. The set of these optimal decision is d^*, whose elements $d_t^*(x_{it}, \eta_i, s_i, \epsilon_{it})$ are:

$$d_t^*(x_{it}, \eta_i, s_i, \epsilon_{it}) = \arg\max_d E_d \left(\Sigma_{k=t}^T \delta^{k-t} \left\{ [u(d_{ik}, x_{ik}, s_i) + \epsilon_{ik}(d_{ik})] | x_{it}, \eta_i, s_i, \epsilon_{it} \right\} \right) \tag{11}$$

where the expectation is taken over the future states induced by d_i^*. The value function at time t, $V(x_{it}, \eta_i, s_i, \epsilon_{it})$, is the expected present discounted value of lifetime utility from following d_i^*, and is given by:

$$V(x_{it}, \eta_i, s_i, \epsilon_{it}) = \max_d E_d \left(\Sigma_{k=t}^T \delta^{k-t} \left\{ [u(d_{ik}, x_{ik}, s_i) + \epsilon_{ik}(d_{ik})] | x_{it}, \eta_i, s_i, \epsilon_{it} \right\} \right) \tag{12}$$

By Bellman's principle of optimality, the value function can also be obtained using the recursion:

$$V(x_{it}, \eta_i, s_i, \epsilon_{it}) = \max_{d_{it}} \left\{ u(d_{it}, x_{it}, s_i) + \epsilon_{it}(d_{it}) + \delta E_d[V(x_{it+1}, \eta_i, s_i, \epsilon_{it+1})|d_{it}, x_{it}, \eta_i, s_i] \right\} \tag{13}$$

Buyers face this dynamic optimization problem – waiting is costly, but may fetch better bids in future. For example, in future the buyer may receive a low-priced bid from a high reputation worker. However, they face uncertainty over the number of bids they will receive in the future and the attributes of those bids.

6   Estimation

We face two key challenges in estimating the model – the size of the state space and the presence of persistent unobserved heterogeneity. In general, high-dimensional state spaces are intractable with nested fixed point algorithms and are estimated using computationally light two-step methods pioneered by Hotz and Miller (1993). However, our state space is intractable even with standard two-step methods. Moreover, two-step methods have traditionally suffered from their inability to account for persistent unobservables (Aguirregabiria and Mira, 2010). While Aguirregabiria and Mira (2007) have proposed a recursive two-step Nested Pseudo Likelihood estimator, it is not applicable here because it requires stationarity and involves large matrix inversions. We therefore adapt the two-step estimation framework recently developed by Arcidiacono and Miller (2011). Their methodology offers two key innovations that we implement:

- First, they generalize Altug and Miller (1998) and provide a framework to exploit finite dependence for a large class of problems, including those with correlated GEV errors that lead to nested logit probabilities (such as ours).
- Second, they present an EM-like algorithm that allows for finite unobserved types within a two-step estimator.

Applications of this framework have been limited. Murphy (2013), Beresteanu et al. (2010), and Ellickson et al. (2012) exploit finite dependence to perform value function reformulations to simplify computation in large state space problems. Finger (2008) and Chung et al. (2010) employ the EM-like algorithm to accommodate persistent unobservables in two-step methods. The latter is a notable early application of this method in the marketing context. In our estimation, we exploit both aspects of this framework.

This estimation strategy has three important benefits in our context. First, given our high-dimensional state space, it makes estimation feasible without having to resort to artificial discretization of the state space or other state-space reduction methods (such as compressing all bid related variables into one inclusive value as in Gowrisankaran and Rysman (2012)). Second, even in our extremely large state space, it allows us to incorporate persistent unobserved heterogeneity.[12] Third, it is robust to mis-specifications in buyers' future expectations far out in the future. For example, in our model, the number of bids that first period dropouts expect to receive in future periods is based on estimates derived from the number of bids received by those who didn't drop out. However, the expectations of these first-period dropouts could be very different. Hence, forecasts of future states are always susceptible to errors because they involve predictions of off-equilibrium paths that are never observed in the data. Because this estimation procedure only projects finite periods into the future (one in our case), it is robust to such mis-specifications. Moreover, because it only simulates one period ahead, it is subject to less simulation bias, even in cases where the state space is sparsely populated. See Arcidiacono and Ellickson (2011) for details.

---

[12]Note that even if we could reduce the state space somewhat, full solution models are still infeasible in our setting because they would require us to embed a standard nested fixed point algorithm inside the EM loop, which is highly computationally expensive. Indeed, EM algorithms are known to be slow even in non-dynamic settings.

6.1 CCP Representation of the Problem using Finite Dependence

We define the ex-ante value function, $V'(x_{it}, \eta_i, s_i)$, as follows:

$$V'(x_{it}, \eta_i, s_i) = \int V(x_{it}, \eta_i, s_i, \epsilon_{it})g(\epsilon_{it})d\epsilon_{it} \tag{14}$$

$V'(x_{it}, \eta_i, s_i)$ is i's continuation value of being in state $\{x_{it}, \eta_i, s_i\}$, after integrating out ϵ_{it}. The continuation value gives us the choice-specific value function, $v(d_{it}, x_{it}, \eta_i, s_i)$, as:

$$
\begin{aligned}
v(d_{it}, x_{it}, \eta_i, s_i) &= u(d_{it}, x_{it}, s_i) + \delta \int V'(x_{it+1}, \eta_i, s_i)f(x_{it+1}|d_{it}, x_{it}, \eta_i, s_i)dx_{it+1} \\
&= u(d_{it}, x_{it}, s_i) + \delta \int V'(x_{it+1}, \eta_i, s_i)f(x_{it+1}|d_{it}, x_{it}, \eta_i)dx_{it+1}
\end{aligned} \tag{15}
$$

$v(d_{it}, x_{it}, \eta_i)$ is i's present discounted value from choosing action d_{it} in period t (net of $\epsilon_{it}(d_{it})$) and following the set of optimal actions d_i^* from period $t + 1$. The second line of Equation (15) follows from the fact that state transitions, $f(\cdot)$, are independent of s_i. Since choosing a bid and canceling are terminal options, there is no continuation value once either of these options are chosen. So:

$$
\begin{aligned}
v(\mathrm{bid}_j, x_{it}, \eta_i, s_i) &= u(\mathrm{bid}_j, x_{it}, s_i) = W_{ib}(x_{it}, s_i) + Y_{ij}(x_{it}) \tag{16} \\
v(2, x_{it}, \eta_i, s_i) &= u(2, x_{it}, s_i) = 0 \tag{17}
\end{aligned}
$$

However, waiting does have a continuation value. So:

$$
\begin{aligned}
v(1, x_{it}, \eta_i, s_i) &= u(1, x_{it}, s_i) + \delta \int V'(x_{it+1}, \eta_i, s_i)f(x_{it+1}|1, x_{it}, \eta_i)dx_{it+1} \tag{18} \\
&= W_{iw}(x_{it}) + \delta \int V'(x_{it+1}, \eta_i, s_i)f(x_{it+1}|1, x_{it}, \eta_i)dx_{it+1}
\end{aligned}
$$

We use these choice-specific value functions to derive the choice probabilities. Rust (1987) showed that, for GEV errors, there exist analytical relationships between choice probabilities and choice-specific value functions, and that these choice probabilities are analogous to the relationships in static discrete choice models. Let $P(1|x_{it}, \eta_i, s_i)$, $P(2|x_{it}, \eta_i, s_i)$, and $P(\mathrm{bid}_j|x_{it}, \eta_i, s_i)$ be the respective probabilities of waiting, canceling, and picking bid j, given state variables $\{x_{it}, \eta_i, s_i\}$. Then the GEV error structure gives the nested logit probabilities:

$$
P(1|x_{it}, \eta_i, s_i) = \frac{e^{[v(1, x_{it}, \eta_i, s_i)]}}{1 + e^{[v(1, x_{it}, \eta_i, s_i)]} + e^{[W_{ib}(x_{it}, s_i) + \sigma I(x_{it})]}} \tag{19}
$$

$$
P(2|x_{it}, \eta_i, s_i) = \frac{1}{1 + e^{[v(1, x_{it}, \eta_i, s_i)]} + e^{[W_{ib}(x_{it}, s_i) + \sigma I(x_{it})]}} \tag{20}
$$

$$
P(\mathrm{bid}_j|x_{it}, \eta_i, s_i) = \frac{e^{[W_{ib}(x_{it}, s_i) + \sigma I(x_{it})]}}{1 + e^{[v(1, x_{it}, \eta_i, s_i)]} + e^{[W_{ib}(x_{it}, s_i) + \sigma I(x_{it})]}} \frac{e^{\frac{Y_{ij}(x_{it})}{\sigma}}}{\sum_{q=3}^{J_{it}} e^{\frac{Y_{iq}(x_{it})}{\sigma}}} \tag{21}
$$

where $I(x_{it}) = \ln \left[\sum_{q=3}^{J_{it}} e^{\frac{Y_{iq}(x_{it})}{\sigma}} \right]$ is the inclusive value of the bid nest. Using the above analytical expressions, we can estimate the model primitives, $\{\alpha, \beta, \delta\}$, if we can compute $\int V'(x_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$ and integrate out the unobserved state variable s_i. Arcidiacono and Miller (2011) show that, for GEV errors, there exists a simple analytical relationship between the continuation value $V'(x_{it}, \eta_i, s_i)$ and the conditional choice probabilities. In our setting, we can express $V'(x_{it}, \eta_i, s_i)$ as:

$$V'(x_{it}, \eta_i, s_i) = \gamma + v(2, x_{it}, s_i) - \ln(p(2|x_{it}, \eta_i, s_i)) = \gamma - \ln(p(2|x_{it}, \eta_i, s_i)) \qquad (22)$$

where γ is the Euler's constant and $p(2|x_{it}, \eta_i, s_i)$ is the conditional probability of choosing the terminal action cancel given observed state variables x_{it} and the unobserved state variables $\{\eta_i, s_i\}$. Because of the finite dependence in the model, the choice-specific value function for cancel has no continuation value, and $\int V'(x_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$ simplifies to:

$$\int V'(x_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1} = \int \left[\gamma - \ln(p(2|x_{it+1}, \eta_i, s_i)) \right] f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1} \qquad (23)$$

In order to evaluate the right hand side of Equation (23), we only need one-period ahead CCPs, $p(2|x_{it+1}, \eta_i, s_i)$, and state transition probabilities, $f(x_{it+1}|1, x_{it}, \eta_i)$.

The choice of 'cancel' as the base option helps in computation because $v(2, x_{it}, s_i) = u(2, x_{it}, s_i) = 0$. So we do not need estimates of the structural parameters associated with $u(2, x_{it}, s_i)$ to obtain numerical estimates of $\int V'(x_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$. On the other hand, if we had chosen 'wait' as the base option, we would need to use updated estimates of $u(2, x_{it}, s_i)$ at each step of the EM to derive $\int V'(x_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$s, which is computationally more intensive.[13]

6.2 Log-likelihood and Estimation Outline

Let π_k be the population probability of a buyer being of unobserved type k, where $k \in \{1, ..., S\}$. Because we have the full history for each auction, we do not condition the initial distribution of types on observed state variables.[14] The joint log-likelihood of i's decision d_{it}s and observed state variables x_{it}s, conditional on x_{it-1}s and unobserved state variables $\{\eta_i, s_i\}$ is given by:

$$L_i(d_i, x_i|\alpha, \beta, \delta, \theta) = \ln \left[\sum_{k=1}^{S} \pi_k \left[\prod_{t=1}^{T_i} \left[\Pr(d_{it}|x_{it}, \eta_i, s_i = k)^{I(d_{it})} \cdot f(x_{it}|d_{it-1}, x_{it-1}, \eta_i) \right] \right] \right] \qquad (24)$$

---

[13] While we have represented $V'(x_{it}, \eta_i, s_i)$ as a function of the terminal action cancel, it can also be represented as a function of the terminal action of choosing a bid, *i.e.*, we can write it as $V'(x_{it}, \eta_i, s_i) = \gamma + v(\text{bid}_j, x_{it}, s_i) - \sigma \ln(p(\text{bid}_j|x_{it}, \eta_i, s_i)) - (1 - \sigma) \ln \left(\sum_{q=3}^{J_{it}} p(\text{bid}_q|x_{it}, \eta_i, s_i) \right)$. However, since $v(\text{bid}_j, x_{it}, s_i)$ and σ are unknown prior to the second stage estimation, this would require us to substitute the updated values for these expressions at each iteration of the EM algorithm, which increases computational time and complexity. In contrast, specifying $V'(x_{it}, \eta_i, s_i)$ as a function of a terminal action that is part of a singleton nest, and whose flow utility has been normalized to zero, simplifies the computation considerably we are now only required to substitute the updated values of CCPs at each step of the EM algorithm.

[14] All buyers start with the same set of observed time varying state variables, *i.e.*, everyone starts at time $t = 1$ and zero bids. However, we could still condition the initial distribution of types on buyer-specific state variables. The results from such an expanded model (available from the author) are not very different from those presented here.

where θ is set of parameters and the non-parametric distributions associated with the state transitions, *i.e.*, $\theta = \{\theta_p, \theta_g, \theta_t, \psi_c\}$, $\Pr(d_{it}|x_{it}, \eta_i, s_i)$s are the choice probabilities for alternatives available to i at period t, and $I(d_{it})$ is an indicator variable that is 1 if d_{it} is the observed choice in the data. Note that we have integrated out the unobserved state variable s_i from the log-likelihood, but not η_i because η_is can be consistently estimated from the data. (See §6.3.1 for details.) Since state transitions are independent of s_i after accounting for η_i, we can write Equation (24) as:

$$L_i(d_i, x_i|\alpha, \beta, \delta, \theta) = \ln\left[\sum_{k=1}^{S} \pi_k \prod_{t=1}^{T_i} \Pr(d_{it}|x_{it}, \eta_i, s_i = k)^{I(d_{it})}\right] + \ln\left[\prod_{t=1}^{T_i} f(x_{it}|d_{it-1}, x_{it-1}, \eta_i)\right] \quad (25)$$

Because the log-likelihoods of state transitions and observed choices are additively separable, they can be maximized separately as follows:

- Estimate state transitions – First, we estimate the Poisson bid arrival process and also obtain consistent estimates, $\hat{\eta}_i$, of η_i. Next, we estimate the non-parametric distributions of bids, the multinomial logit model of seller's geographic region, and the logit model of buyer-seller past interaction. Since these models don't depend on s_i, they can be consistently estimated at the first stage.

- Augmented two-step EM estimator – Recursively compute and update the CCPs, π_ks, structural parameters $\{\alpha, \beta\}$, and discount factor δ. At this stage, estimates from the state transition models are used to calculate the future continuation values in conjunction with CCPs.

Without persistent unobserved heterogeneity s_i, we could have used a standard two-step estimator, by estimating CCPs along with the state transition models, as functions of x_{it} and $\hat{\eta}_i$. (Because η_i doesn't affect flow utilities, a consistent estimate, $\hat{\eta}_i$, of η_i is available after the estimation of the Poisson bid arrival process.) However, the inclusion of s_i necessitates the use of EM algorithm, since we can no longer obtain consistent estimates of CCPs from the data (s_i and π_ks are unknown).

Estimating the model in stages does not affect the consistency of the results (Rust and Phelan, 1997; Rothwell and Rust, 1997), though it does lead to lower standard errors for structural parameters because estimates of state transition probabilities and CCPs are treated as data in the estimation. Hence, we bootstrap these standard errors.[15] We now describe each estimation step in detail.

6.3 Estimation of State Transitions: Models of Bid Arrival and Bid Attributes

6.3.1 Estimation – Fixed Effects Poisson

The set of parameters to be estimated is $\theta_1 = \{\theta_p, \eta_1, ..., \eta_N\}$, which includes the fixed effects (η_is). We use the maximum likelihood approach to estimate θ_1. Below, we present an overview of the estimation and refer interested readers to Winkelmann (2003) for details.

Let b_i be the vector of number of bids received by i, where $b_i = (b_{i1}, b_{i2}, ..., b_{iT_i})$, and T_i is the last period in which the auction is still active. For example, if i cancels her auction two days after posting it,

---

[15]We sample on auctions and use 250 replications in our bootstrap procedure.

then $T_i = 2$. Equation (26) gives us the conditional log-likelihood contribution of i as:

$$L_{ip}(b_{i1}, ..., b_{iT_i}|\theta_p, \eta_i) = -\eta_i \sum_{t=1}^{T_i} \lambda_t + \ln \eta_i \sum_{t=1}^{T_i} b_{it} + \sum_{t=1}^{T_i} b_{it} \ln \lambda_t - \sum_{t=1}^{T_i} (\ln b_{it})! \qquad (26)$$

Setting the first derivative of $L_{ip}(b_{i1}, ..., b_{iT_i}|\theta_p, \eta_i)$ to zero, we have $\hat{\eta}_i = \frac{\Sigma_{t=1}^{T_i} b_{it}}{\Sigma_{t=1}^{T_i} \lambda_t}$. Substituting this back into Equation (26) gives us a conditional log-likelihood, $L_{ip}(b_{i1}, ..., b_{iT_i}|\theta_p)$, which is independent of η_i. ($\Sigma_{t=1}^{T_i} b_{it}$ is a sufficient statistic for η_i.) Then, the conditional log-likelihood of all the bid arrivals observed in the data is given by:

$$L_p(\theta_p) = \sum_{i=1}^{N} L_{ip}(b_{i1}, ..., b_{iT_i}|\theta_p, \hat{\eta}_i) \qquad (27)$$

We then maximize $L_p(\theta_p)$ to estimate θ_p. Once we have a consistent estimate, $\hat{\theta}_p$, of θ_p, we can use it to consistently estimate η_i using: $\hat{\eta}_i = \Sigma_{t=1}^{T_i} b_{it} / \Sigma_{t=1}^{T_i} \hat{\lambda}_t$, where $\hat{\lambda}_t = \exp(z_t' \hat{\theta}_p)$.

6.3.2 Estimation – Non-Parametric Joint Distributions

To obtain the non-parametric joint distributions, we first estimate the optimal bandwidth $h_c \forall c \in \mathcal{C}$ using likelihood cross-validation (Duin, 1976; Silverman, 1986). Let $\hat{\psi}_c(h, A)$ and $\hat{\psi}_{c,-q}(h, A)$ be the p.d.f estimate of point A from the c^{th} category using bandwidth h and datasets $\{A_1, ..., A_{M_c}\}$ and $\{A_1, ..., A_{q-1}, A_{q+1}, ..., A_{M_c}\}$, respectively. Then the cross-validation score of h for category c is given by averaging the log-likelihood $\hat{\psi}_{c,-q}(h, A)$ over all q.

$$CV_c(h) = M_c^{-1} \sum_{q=1}^{M_c} \ln[\hat{\psi}_{c,-q}(h, A)] \qquad (28)$$

The likelihood cross-validation choice of the optimal bandwidth is the value that maximizes $CV_c(h)$. Intuitively, the cross-validation score $CV_c(h)$ is the log-likelihood of observing the dataset. $\hat{\psi}_{c,-q}(h, A)$ is the probability of drawing the data-point A (assuming it is not part of the dataset). So $M_c^{-1} \Pi_{q=1}^{M_c} \hat{\psi}_{c,-q}(h, A)$ is the total probability of observing the dataset.

While the maximization is conceptually simple, it is computationally intensive. To evaluate the likelihood at a given bandwidth, we need to evaluate the density at each data point at that bandwidth and then sum over the density contributions of all data points. Moreover, at each data point, we need to find the k^{th} nearest point in the Euclidean space to calculate its density contribution. This becomes prohibitively expensive as the size of the dataset and the number of dimensions increase. Moreover, this has to be done multiple times to reach the optimal h_c. While many algorithms have been proposed to address these computational issues, we follow the recent method proposed by Gray and Moore (2003), which is based on *k-d trees* and has been shown to be much faster than previous methods. We use the MATLAB-based KDE toolbox to perform the estimation (Ihler, 2003).

6.3.3 Estimation – Logit Models of Seller Region and Past Buyer-Seller Interaction

The estimation of the multinomial logit model of seller region (bg_{ij}) and the binary logit model of past buyer-seller interaction indicator (bt_{ij}) are straightforward. The log-likelihood of drawing the sellers' geographic

regions observed in the data is:

$$L_g(\theta_g) = \sum_{i=1}^{N} \sum_{j=3}^{J_{iT_i}} \sum_{q=1}^{4} \ln[h_g(bg_{ij} = q|gx_{ij}, \theta_g)^{I(bg_{ij}=q)}] \tag{29}$$

where $J_{iT_i} - 2$ is the total number of bids that i has accumulated in the last period that she is active (T_i). Similarly, the log-likelihood of the buyer-seller interactions observed in the data is given by:

$$L_b(\theta_t) = \sum_{i=1}^{N} \sum_{j=3}^{J_{iT_i}} \sum_{q=1}^{2} \ln[h_b(bt_{ij} = q|tx_{ij}, \theta_t)^{I(bt_{ij}=q)}] \tag{30}$$

Maximizing the above log-likelihoods gives us consistent estimates of θ_g and θ_t.

6.4 Two-step EM Estimator

In the second stage, we estimate the CCPs, the population probabilities of unobserved types (π_ks), and the structural parameters and discount factor $\{\alpha, \beta, \delta\}$. The first part of the log-likelihood from Equation (25) suggests the maximization:

$$(\hat{\alpha}, \hat{\beta}, \hat{\delta}, \hat{\pi}) = \arg\max_{\alpha,\beta,\delta,\pi} \sum_{n=1}^{N} \ln \left[\sum_{k=1}^{S} \pi_k \prod_{t=1}^{T_i} \Pr(d_{it}|x_{it}, \eta_i, s_i = k)^{I(d_{it})} \right] \tag{31}$$

Dempster et al. (1977) note that the first-order condition for the above maximization problem is same as that of the following maximization if $\rho(k|d_i, x_i; \hat{\alpha}, \hat{\beta}, \hat{\delta}, \hat{\pi})$ were treated as known:

$$(\hat{\alpha}, \hat{\beta}, \hat{\delta}) = \arg\max_{\alpha,\beta,\delta} \sum_{n=1}^{N} \prod_{t=1}^{T_i} \sum_{k=1}^{S} \rho(k|d_i, x_i; \hat{\alpha}, \hat{\beta}, \hat{\delta}, \hat{\pi}) \cdot \ln \left[\Pr(d_{it}|x_{it}, \eta_i, s_i = k)^{I(d_{it})} \right] \tag{32}$$

$\rho(k|d_i, x_i; \hat{\alpha}, \hat{\beta}, \hat{\delta}, \hat{\pi})$ is the posterior probability of i belonging to unobserved type k given data (d_i, x_i), population type estimates $\hat{\pi}$, and structural parameter estimates $\{\hat{\alpha}, \hat{\beta}, \hat{\delta}\}$. Since ρ is unknown, Dempster et al. (1977) proposed a recursive EM algorithm that starts with a set of assumed structural parameters, based on which ρs are updated, which in turn are substituted back into the maximization problem in Equation (32) to obtain a new set of parameters. This process is repeated till the parameters and πs converge.

In a standard finite-mixture setting, this is relatively straightforward to implement. However, in a dynamic setting, the choice probabilities, $\Pr(d_{it}|x_{it}, \eta_i, s_i)$, are functions of the unknown continuation values through $\int V'(x_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$, which have to be calculated using the analytical expression shown in Equation (23). In order to do so, we first need estimates of the CCP of cancellation, $p(2|x_{it+1}, \eta_i, s_i)$, which cannot be directly obtained from the data since s_is are unknown. Arcidiacono and Miller (2011) propose an expanded version of the EM algorithm where the CCPs are also updated at each step of the EM algorithm. We follow their approach. Appendix 1 provides the step-by-step details of the estimation process.

6.5 Identification

6.5.1 Identification of CCPs, State Transitions, Population Distribution of Types

Non-parametric identification of CCPs, state transitions, and the population distribution of types is an important prerequisite for using the two-step methods. While persistent unobserved heterogeneity can lead to non-identification, Kasahara and Shimotsu (2009) and Arcidiacono and Miller (2011) show that in many cases identification can be restored. Specifically, they prove that if all states are reachable after all decisions (*i.e.*, unconstrained state-space evolution), then the number of decision-state sequences available for identification expands exponentially with time and state-space size. In this context, they show that data on three or more timeperiods is sufficient for non-parametrically identifying CCPs and population probabilities.

However, in our setting, all states are not reachable after all decisions. Waiting is the only continuation decision in our model and the only decision after which the state space can change. Moreover, once a buyer has received a certain number of bids, she cannot go back to fewer bids or change the attributes of the bids she has already received. This constrains state space evolution and restricts the number of decision-state sequences that are observable in the data. So the general proofs by Kasahara and Shimotsu (2009) and Arcidiacono and Miller (2011) are not directly applicable to our setting.[16] However, we do have a large number of time periods, the number of bids that a buyer can receive is high, and for a given number of bids, the possible combinations of bid attributes is very large. Therefore, we still have sufficient observed decision-state sequences to identify CCPs and state transitions.

In Appendix 2 we present a formal proof. Here, we provide some intuition on how our CCPs are identified. Since state space advances only through wait decisions, by default, CCPs of waiting appear in many decision-state sequences and are easily identified; *i.e.*, a buyer in a given state could have jumped into it in many ways, and with the wait decision, can also jump out of it in many ways. However, termination CCPs (CCPs of canceling and choosing a bid) appear only in sequences that end in the current state and action. Hence, one might suspect that they are not identified. However, that is not the case because the same terminal decision-state combination can be reached in many ways. For example, the CCP of canceling for a buyer of a specific unobserved type with 3 bids in period 2 will appear in all of the following sequences – (a) buyer receives 0 bids in $t = 1$, waits, and receives 3 in $t = 2$, (b) receives 1 in $t = 1$, waits, and receives 2 in $t = 2$, (c) receives 2 in $t = 1$, waits, and receives 1 in $t = 2$, and (d) receives 3 in $t = 1$, waits, and receives none in $t = 2$. Note that while all these sequences end with the same state and terminal decision, there are still enough sequences for identification. Of course, with bid attributes, the number of such sequences increases even more. This allows terminal CCPs to be identified.

Nevertheless, we don't have sufficient degrees of freedom to nonparametrically identify more than two

---

[16]We present a small example to demonstrate the differences in the identification of problems with constrained state-space evolution and those without constraints. Consider a general case with ten possible states and five possible decisions that the agent can make at any point in time. Let each decision be a continuation decision such that the agent can transition into any of the states in the next period following any decision. Then, with T periods, the number of observable decision-state sequences is $(10 \times 5)^T$. If the number of CCPs and state transitions that we need to identify is less than this number, then we have ensured that the first-stage estimates are non-parametrically identified. Now consider a modification of the previous case so that only one decision is a continuation decision, and the rest are terminal. Then the number of decision state sequences is $\Sigma_{t=1}^{T} 10^t$. This is considerably lower than $(10 \times 5)^T$. Hence, constraints on the state space evolution can interfere with the non-parametric identification of first stage estimates and should be verified for each application setting.

types in this context. We formally show why this is the case in §2 in the Web Appendix. Thus, in our application, we only allow for two unobserved types.

6.5.2 Identification of Discount Rate and Utility Parameters

The discount factor is identified through exclusion restrictions (Rust, 1994; Magnac and Thesmar, 2002). The most important restriction we employ is the exclusion of time from flow utility, *i.e.*, time affects state transitions through bid arrival rates, but not flow utilities. Apart from time, buyer-auction fixed effects (η_is) also help in identifying the discount factor since η_is affect bid arrival rates but are not included in flow utilities. (The first stage estimate, $\hat{\eta}_i$, of η_i is treated as known in the second stage.) Note that part of the variation in η_is is explained by buyer and auction specific variables that are also included in the flow utility (*e.g.*, buyer reputation, MaxBid, unobserved auction type s_i), and hence this variation cannot contribute to the identification of the discount factor. However, we find that even after accounting for observed time-invariant buyer, auction specific state variables and s_i, a significant amount of variation in η_i still remains unexplained; and it is this remnant variation that acts as an exclusion restriction (i.e., it affects bid arrivals, but not flow utility) and helps in identifying the discount factor.

Utility parameters associated with observed buyer and bid attributes are identified as usual – from the variation in data and buyers' choices – so we don't go into their details. Instead we focus on unobserved types (s_i), which are identified through the dynamics of the model, realizations of bid arrivals, bid attributes, and buyers' decisions. For example, a buyer with High unobserved preference for bids may get only one bid in period 1 and pick it right away, while a buyer with low unobserved preference for bids may get many bids over time and repeatedly wait. Consider two buyers who are identical on $(x_{it}, \hat{\eta}_i)$ and receive similar bids in the first period. Suppose one picks a bid at $t = 1$ and the other continues for two weeks without picking bids. Without unobserved heterogeneity, both these buyers would have the same predicted bid choice probabilities. However, with unobserved heterogeneity, we learn more about both buyers – the first buyer has a clear preference for choosing bids, either because the bids she receives are of high unobserved quality, or because she has no good outside options, while the second buyer clearly has low inherent preference for choosing bids, either because her bids are of low unobserved quality or because she has good outside options. These kinds of variations in the data help identify types.[17]

7 Results

7.1 State Transition Estimates: Bid Arrival and Bid Attributes

We now discuss a few highlights from the first step results here and refer interested readers to the accompanying Web Appendix for details on the parameter estimates.

First, in the context of the Poisson model, our estimates suggest that bid arrivals slow down considerably with time. We also find that there is significant variation in the estimated fixed effects, $\hat{\eta}_i$s. This affirms the importance of including auction specific unobserved heterogeneity (η_i) in our bid arrival model. Overall,

---

[17]Note that allowing for unobserved heterogeneity in flow utilities across buyers through s_i and holding discount rate constant across buyers is analogous to allowing for unobserved heterogeneity in discount rates and keeping flow utilities independent of persistent unobservables. So buyers who exhibit an aversion to picking bids (type $s_i = 1$, see Table 8) can be interpreted as patient buyers and vice-versa. In other words, unobserved heterogeneity in buyers' patience and preference for bids cannot be separately identified in our setting.

we find that the Poisson model captures the patterns very well. Second, from the multinomial logit model of seller region, we find that bid price and many of the bidder attributes are correlated with the bidder's geographic region. A seller's geographic region is also influenced by the geographic region of the buyer and her reputation. For example, everything else being constant, buyers are more likely to attract sellers from their own region. This effect likely stems from lower communication costs and similarities in IP restrictions within a region. Third, the estimates from the logit model for the indicator of past buyer-seller interaction suggest that only a small percentage of sellers are likely to have interacted with the buyer. However, the probability of drawing such sellers is higher for buyers who specify higher MaxBids, have initiated a large number of uncanceled auctions in the past, and have a good reputation. Similarly, sellers who belong to regions 1 and 2, quote slightly higher prices than average, and have a good reputation on the site are more likely to have interacted with the buyer in the past.

Finally, we discuss our estimates of the non-parametric joint distributions of bid price, number of bidder ratings, and bidder average rating. There are no parametric results in this context except the bandwidths (h_c) for the 40 sub-classes. Since these bandwidths are not very informative in and of themselves, we do not present them here. However, we note that the kernel density estimates are, in general, very good at approximating the joint distributions of these bid attributes; Koglomorov-Smirnov tests comparing the equality of the original and estimated distributions for all of the sub-classes confirm them to be indistinguishable.

7.2 CCP Estimates

Due to the size of the state space and the presence of a large number of continuous state variables, CCPs are estimated using flexible logits. We include all of the state variables, their higher order terms and interactions in our flexible logit. Use of flexible logits to model CCPs has precedence in the literature (Arcidiacono and Miller, 2011; Fang and Wang, 2013; Murphy, 2013). Because of the large number of CCP parameters and the difficulty in interpreting them, we do not present the CCP estimates here. Instead, we highlight the main effects. A buyer's probability of canceling an auction increases concavely with the prices of the bids received. It however decreases with the mean ratings of those bidders, the MaxBid of the auction, the number of bids received so far, and the number of past ratings received by the buyer. The probability of cancellation also varies with unobserved buyer type and time.

7.3 Structural Parameter Estimates

The estimates of the structural parameters associated with buyers' utility are presented in Table 8. The first column (Model M1) presents a static model that ignores both dynamics and persistent unobservables. The second column (Model M2) presents a model with dynamics but without the unobserved state variable s_i. Here, we do not control for persistent unobserved heterogeneity in buyers' taste for bids or the average quality of bids received by a buyer. In the third column, we consider a model (Model 3) with two unobserved types $s_i \in \{1, 2\}$. Unless specified otherwise, we discuss the results from Model 3 throughout.

7.3.1 Bottom Level Utility Parameters

The top set of rows show the estimates for coefficients that vary within the bid nest. First, we find that buyers' utility from bids is decreasing in price, though concavely, *i.e.*, the same price increase is less painful at higher prices. Next, we find that sellers with more ratings are more attractive, *i.e.*, the coefficient of $\ln(bn_{ij} + 1)$

is positive. This suggests that buyers value the site's reputation system – they prefer sellers who have been in the system for some time and on whom information, in the form of past ratings, is available. Note that this positive effect is consistent with the cheap-identity problem prevalent in Internet marketplaces. Sellers with few past ratings are not only untested (and therefore less trustworthy), but also more likely to be low quality sellers, who have simply reappeared as a new seller after milking their old reputations. However, the marginal value of each additional rating is decreasing, possibly because the new information in each new rating decreases as the number of ratings increase. This effect is consistent with 'imperfect monitoring', because the number of reviews would be irrelevant if reviews were perfectly informative. (We also tried other functions of number of ratings, and found the fit of the model with $\ln(bn_{ij} + 1)$ to be the best.)

Buyers also derive value from the mean rating of the seller (*i.e.*, the coefficient of $(br_{ij} - \bar{br})$ is positive). If the seller's mean rating is higher than of the average seller in the marketplace, then the buyer derives a positive value, else the buyer derives a negative value. Note that this effect suggests that buyers adjust for the 'rating inflation' problem, commonly observed in online reputation systems. The square term of centered mean rating term also has a positive effect. That is, any increase in a seller's mean rating above the mean ($\bar{br}$) has a convex benefit, while the opposite is true for decreases below the mean rating. Moreover, the interaction of the ln of number of ratings and the centered mean rating, $(br_{ij} - \bar{br})\ln(bn_{ij} + 1)$, is also positive. This implies that buyers' valuation of a marginal change in a seller's rating (over the population mean) is proportional to the ln of the number of ratings she has received. That is, a seller with a mean rating of 9.5 and 20 ratings is valued less than a seller of mean rating 9.5 and 25 ratings, even after controlling for the main effects of the number of ratings. Again, this effect reflects the fact that ratings are noisy.

MaxBid has no impact on a buyer's valuation for seller reputation. If we interpret MaxBid as project size or buyer's value for the project, then we can conclude that buyers don't necessarily place more value on seller reputations for larger or more important projects. A buyer's own reputation also doesn't influence her valuation of sellers' reputations. Buyers have a strong preference for sellers with whom they have worked in the past. In principle, the direction of this effect could go either way, depending on whether the previous interaction went well or not. However, in the data, we find that, usually sellers avoid buyers who gave them low ratings previously. So the indicator for previous interaction almost always indicates a 'good' previous interaction. The strong positive effect is then understandable since the information asymmetry problem is alleviated when the buyer has traded with the seller in the past and liked her.

Finally, we find that buyers prefer sellers from developed countries the most, followed by those from eastern European countries, and buyers in developing countries have a small preference for sellers from their own geographic regions, possibly because such sellers might be easier to communicate with and/or have a better understanding of the local intellectual property regulations.

7.3.2 Top Level Utility Parameters

The two bottom sets of rows show the top level coefficients for the wait and bid nests. These coefficients represent the relative attractiveness of the two nests, in comparison to canceling.

We find that buyers' flow utility from waiting is increasing in MaxBid and decreasing in the number of buyer ratings. Buyers who are new to the site, those who have few ratings, those who have few uncanceled auctions, and those with high cancel ratios receive lower utility from picking a bid. This is possibly because

buyers with a good reputation and history in this specific freelance community have higher costs of finding alternative workers from outside sources, after canceling. Further, buyers from region 1 are less likely, and those from region 2 are more likely to choose a bid. This might be due to local labor costs, which are high in region 2 and low in region 1.

We also find that the High types ($s_i = 2$) have high unobserved preference for bids compared to the Low types ($s_i = 1$). The two types are distributed in the ratio of about 3:7, with the High type being in the minority.

7.3.3 Nest Correlation and Discount Factor

The last two rows show the estimates of the bid nest's correlation parameter σ and the discount factor δ. The nesting parameter σ in Model 3 is 0.4553, significantly less than one, which suggests that the unobservable preferences for bids, ϵ_{ijt}, have a component that is correlated across the bid options. Moreover, note that the addition of persistent unobserved heterogeneity in buyers' taste for bids, in the form of s_i, reduces the contemporaneous correlation between the ϵ_{ijt}s. This suggests that σ was picking up some of the persistence in buyers' preference for bids in Model 2.

The model without persistent unobserved heterogeneity in preferences (Model 2) attributes the early exit of a large fraction of buyers to impatience and significantly underestimates the discount factor. This underestimation is rectified in Model 3, where the daily discount factor is estimated to be 0.8823. While this seems reasonable for this context, it is lower than the discount rate implied by the yearly interest rate. Given the size of these jobs (at most worth $100), this level of impatience is understandable – it is relatively easy for a buyer to complete the job herself in a few hours.[18] Our findings highlight the importance of empirically estimating the discount factor, especially in less conventional non-monetary settings, where traditional discount factors calculated from interest rates are unlikely to be applicable. Please see Frederick et al. (2002) for an exhaustive review of experimental studies on time discounting. Dube et al. (2011), Chung et al. (2010), and Yao et al. (2011) also provide excellent discussions on the magnitude of estimated discount factors using field and survey data.

8 Validation

We now compare the fit and performance of our model with two inferior models – a static model and a dynamic model without persistent unobserved heterogeneity.

In the static model, only the final outcomes are analyzed – cancel or choose a submitted bid (Model M1 in Table 8). We find that the static model considerably overpredicts cancellation for early deciders (by 49.92% for cohort E1) and underpredicts it for late deciders (by 44.9% for cohort E13). See Table 9 for details. As mentioned in §1, there are two issues here – dynamics and self-selection. First, buyers who choose to exit early are likely to have drawn a very good set of bids, making them less likely to cancel. In contrast, those who wait are likely to have drawn a poor set of bids, and therefore more likely to cancel. Second, buyers who wait longer are a self-selected group – their repeated decisions to not choose a bid indicate that they either have good unobserved outside options and/or an inherently low taste for bids.

---

[18]As a robustness check, we also estimated a version of the model with the discount factor fixed at 1 and confirmed that reputation effects are qualitatively similar to those presented here.

Hence, when they do exit, they are more likely to cancel. Because the static model cannot account for these factors, it does a poor job of predicting the realized outcomes in the data.

Next, in Table 10, we present the fit of two dynamic models, with and without persistent unobserved heterogeneity (Models 2 and 3 from Table 8). Actual and model predicted probabilities of canceling and choosing a bid for surviving buyers for each period are shown. To calculate the predicted probabilities, we first calculate the continuation values using our estimates of the CCPs and state transitions. Then we plug them and our utility estimates into the closed-form expressions of nested choice probabilities. For the model with unobserved heterogeneity, we weigh the choice probabilities of a buyer (or auction) i with her ex post probabilities of belonging to the two unobserved types conditional on the observed outcomes in data.

While both models perform better than the static model, the model without persistent unobserved heterogeneity is still very poor. Because it doesn't recognize that early bid choosers likely have a high unobserved preference for bids and late exiters likely have low unobserved preference for bids, it underpredicts bid choice initially and overpredicts bid choice in later periods. The underprediction in the first period is especially severe, when the observed choice probability for the bid nest is 0.4921 and the model predicted probability is 0.3213. While it performs better in predicting cancellations, the predictions worsen with time. For example, at $t = 12$, the observed probability of cancellation is 0.1813, while the model predicted probability is 0.1032. Overall, these fit values suggest that a dynamic model alone is not sufficient to explain the realized outcomes in the data.

Finally, note that the fit of the dynamic model with persistent unobserved heterogeneity is very good, especially when compared to previous models. It slightly underpredicts bid choice in the first period and overpredicts bid choice in the middle periods. However, it generally doesn't deviate from the observed choice probabilities by large amounts. The improvement in the fit of this model, highlights the importance of controlling for dynamic selection in this context.

9 Simulation Results

Given that the model predicts the realized outcomes reasonably well, we now use our estimates to derive additional results by running simulations. In this section, we run simulations without making any regime changes, and consider regime changes in §10.

9.1 Sellers' Perspective: Returns to Reputation

Below, we present two types of results to highlight the returns to reputation for a seller.

9.1.1 Aggregate Choice Probability Maps

We simulate the expected probability of being successfully picked for a focal bidder, with a certain set of reputation attributes (number of ratings and average rating) and bid price, if she were to be the first bidder in a randomly chosen auction. The variation in the winning probabilities across different reputation values and bid prices are used to calculate the returns to reputation in this market. Simulation details are described in Appendix 2.1 and the results are presented in Figure 7. We have chosen to report the results for MaxBid = 100 because it is the dominant category, though results for other MaxBids are similar.

In Figures 7a and 7b, we set the bid prices at 50 and 100 respectively, and vary the average rating and number of ratings. It is clear that increases in both the average rating and the number of ratings result in

higher winning probabilities. More importantly, the positive interaction effect between number of ratings and the average rating leads to greater returns at higher values of these metrics. For example, at a bid price of $50 and an average rating of 9, going from 1 to 128 ratings improves the probability of success by 0.14, which translates to a 35% increase in expected revenue. On the other hand, the same increase in number of ratings at an average rating of 10 improves expected revenue by about 50%.

In Figures 7c and 7d, we set the number of ratings at 4 and 64 respectively, and vary the bid price and average rating. With four ratings, a seller bidding $100 can go from an expected probability of success of 0.25 at an average rating of 8 to a probability of 0.37 at an average rating of 10, which translates to almost 50% increase in expected revenue. With 64 ratings, this increases to approximately 100%. The percentage increase in revenues at lower bid prices are also significant though lower. For example, at a bid of $40 and 64 ratings, going from an average rating of 9 to 10 increases expected revenues by 20%.

In Figures 7e and 7f, we set the average rating at 9 and 10 respectively, and vary the bid price and number of ratings. At an average rating of 9, going from 1 to 128 ratings increases expected probability of winning by approximately 0.1, which translates to anywhere between a 20-40% increase in expected revenue depending on the bid price. At an average rating of 10, the expected increase in revenue is higher at 40-70%. Overall, we find that there are significant returns to reputation in this marketplace.

9.1.2 Iso-classes of Sellers

We now present *iso-classes* of sellers based on reputation and price. To derive these iso-classes, we pick a focal bidder and then vary the price and reputation metrics while holding the expected utility from the seller constant (using Equation 6). This gives us a set of sellers who provide the same expected utility, and have the same probability of being chosen as the initially chosen focal seller, within a given choice set.

Consider a focal bidder with median reputation who quotes $50. Figure 8a shows the heat-map of prices that bidders with different reputation metrics would have to quote in order to be in the same iso-class as the focal bidder. For example, a seller with a mean rating of 10 and 6 ratings, who quotes $70 is in the same iso-class as the focal seller, which translates to a $20 premium for a 0.66 increase in mean rating. Similarly, Figures 8b and 8c show the iso-classes for median reputation sellers quoting $80 and $100, respectively. Figure 8d shows the iso-class of sellers who have past history with the buyer, when the focal seller quotes $25, is at the 75[th] percentile of reputation, and has no past history with the buyer. Note that even with high reputations and low prices, it can be difficult to compete with sellers who have interacted with the buyer in the past. For instance, a seller with the same reputation as the focal seller but past interaction history, is able to charge $80 and be in the same iso-class.

9.2 Buyer Entry

We now present some results on buyers' entry decision. Even though we did not explicitly model buyer entry, we are able to examine entry decisions because of the structural nature of our model.

Buyer i chooses to post an auction/enter the market at time $t = 0$ if his expected utility from doing so is greater than that from not entering. If we normalize the utility of not posting an auction to zero (similar to that from canceling), and assume that buyers' costs of making the actual post is negligible, we can write out

i's entry decision in period $t = 0$ as:

$$\delta \int V'(x_{i1}, \eta_i, s_i) f(x_{i1} | enter, x_{i0}, \eta_i) + \epsilon_{i0}^{enter} > \epsilon_{i0}^{no-enter} \tag{33}$$

where the first term is the discounted future value of entering the auction (and making optimal decisions henceforth) and right hand side is the utility from not entering. The two error terms, $\epsilon_{i0}^{enter}, \epsilon_{i0}^{no-enter}$ are assumed to be i.i.d extreme value. This gives us the entry probability of buyer i as:

$$P(enter | x_{i0}, \eta_i, s_i) = \frac{e^{\delta \int V'(x_{i1}, \eta_i, s_i) f(x_{i1} | enter, x_{i0}, \eta_i)}}{1 + e^{\delta \int V'(x_{i1}, \eta_i, s_i) f(x_{i1} | enter, x_{i0}, \eta_i)}} \tag{34}$$

We can thus calculate the a priori expected probability of buyer i's entry using Equation (34).[19] In Figure 9, we present the average entry probabilities for the two types of buyers for all MaxBids. There are two points of note here. First, we find that buyers' likelihood of entry is increasing with MaxBid, *i.e.*, buyers are more likely to enter higher value auctions. Second, 92.9% of High type buyers who consider entry in any given period actually enter, whereas only 74.6% of Low type buyers do. This discrepancy stems from the differences in the value that the two types place on bids. Recall that High type buyers have significantly higher value from choosing bids; consequently, they are more likely to enter the auction in anticipation of this future utility. This difference (almost 20%) further highlights the importance of accounting for persistent unobserved heterogeneity in buyers' differences.

10 Counterfactuals

We now consider a set of regime changes and document their impact on auction success rates and firm revenues. Specifically, we seek to answer the following questions:

- What is the value of the site's reputation system?
- Can the site improve its revenues by attracting more sellers?
- Can the site improve its revenues by ridding itself of low reputation, low price sellers and instead attracting high reputation, high price sellers?
- Can the site benefit from charging a fee to post auctions? If yes, should it charge a fixed fee for all auctions or should it instead consider a fee based on the % of the MaxBid?

Note that these questions don't have a priori obvious answers, and hence require an empirical structural model to arrive at reasonable answers. For example, consider the issue of auction fees. On the one hand, auction fees are a direct source of revenue for the firm. On the other hand, they have a negative impact on buyer entry and some buyers who may have previously procured from the site may not even enter, thereby leading to a loss in revenue from commissions. Using our structural framework, we are able to empirically evaluate the impact of such opposing forces and make normative recommendations to the firm.

In each counterfactual, we simulate the auctions and re-solve the buyer's decisions under the new regime. CCPs from the original solution are not valid under regime changes, and we use the full solution method to

---

[19]First period CCPs are not identified, so we use the full solution method to obtain value functions at $t = 1$. While the size of the state space makes this computationally intensive, it remains feasible because we need only to solve for the value functions once using a backward solution. See §4 in the Technical Appendix for details.

obtain the value functions in our simulations. Because each counterfactual only requires us to solve for the value functions once (by specifying the continuation values as functions of inclusive values and employing some amount of discretization), we are able to make the problem computationally feasible. Further, since this is a non-stationary dynamic programming problem, we use the backward solution method to solve for value functions, the details of which are given in §4 of the Technical Appendix. Finally, we also note that because we have a partial equilibrium model, the usual caveats apply when interpreting our results.

10.1 Value of Reputation System

Earlier, we saw that buyers value high reputation sellers – they are not only more likely to pick them, but are also willing to pay higher prices to them. Since the freelance site generate revenues through percentage commissions on prices, a robust reputation system that sustains high equilibrium prices can be a significant source of revenue and competitive advantage to freelance sites. Therefore, using our estimates on the primitives of buyer utility, we now quantify the value of the reputation system for the firm. We keep the distribution of sellers and buyers the same and set the average rating of all sellers to the population mean and the number of ratings to zero. Since we use mean-centered average ratings and ln(number of ratings+1) in buyers' utility model, this sets both the reputation effects to zero. We also assume that the bid arrival process remains the same. Then we re-solve for the buyers' decisions.

In the absence of a reputation system, buyers have lower value from choosing bids, and more of them prefer to cancel the auction. This has a direct negative effect on the site's revenues. Further, since the reputation attributes have now disappeared, buyers' relative weight on price increases. Thus successful auctions now clear at lower prices, which has an additional negative effect on the site's revenues though decreased commissions. Thus, the cumulative impact on revenues is negative and higher than that implied by the lower clearance rates. Overall, we find that revenues fall by 11.1%. Further, we find that the reputation system is more valuable for high value auctions (see Figure 10). Specifically, we find that the revenue loss for auctions with $10 MaxBid is about 5.37%, whereas it more than doubles to 12.51% for $100 auctions. This suggests that the site may benefit from further investments in its reputation system, especially for high value auctions.

Finally, a necessary caveat here is that our assumption that bid arrivals, distribution of sellers, and prices, remains same may not be reasonable. For instance, without the reputation system, high quality sellers may leave the marketplace, leaving only low quality ones, thereby increasing buyers' likelihood of canceling the auction even more and leading to even lower revenues. Thus, the current findings can be interpreted as a lower bound on the inferred value of the reputation system. Further, if the manager has external information (from surveys or previous regime changes) on how sellers modify their behavior in response to changes in the reputation system, they may include it in counterfactual simulations, which can then be used to inform managerial initiatives. Thus, our model and estimation framework can be used as a managerial tool to test the impact of policy changes before implementing them.

10.2 Modifications to Sellers' Side

Next, we present two counterfactual experiments where we modify the supply (seller) side and examine the impact of the regime change on both cancellation rates and revenues.

10.2.1 Increasing the Supply of Sellers

In this simulation, we examine the impact of increasing the supply of sellers without changing the distribution of seller reputations and prices. The firm can achieve this by reducing the commission rates, which would attract new sellers as well as incentivize existing sellers to bid more. The key question that we seek to answer here is – if the firm could engineer each auction to have more bids (without changing the distribution of bids that it receives), would it increase the number of successful transactions and revenues?

We simulate all the auctions in the data a large number of times with the following modification – for each auction, the bid arrival rate is increased by an inflation factor.[20] We track the success (bid choice or cancel, when the auction is terminated) and the transaction price (if bid was chosen) for each simulated auction. The simulation results are presented in Figure 11. The number of successful auctions increases by $\approx 2.5\%$ when the bid arrival rate doubles is inflated by 1.5, and by $\approx 4\%$ when it doubles. This is understandable because as the number of bids increase, a buyer is more likely to draw a high value bid (H-reputation, L-price, *etc.*). Surprisingly though, the increase in the transactions doesn't make much of a dent on the revenues, which increase by a small percentage initially ($\approx 1\%$ at an inflation factor of 1.6), and then start decreasing! This is in spite of the monotonic increase in the number of successful transactions. The lukewarm response of revenues stems from the competition effect – when buyers have a larger number of bidders to choose from, they are more likely to find sellers quoting lower prices (and reasonable reputation values), which drives down the transaction price.[21] Since the firm's revenues come from percentage commissions on the transaction price, lower transaction prices lead to lower revenues.[22] Hence, our results suggest that just increasing the supply of sellers on the site is not sufficient to increase revenues.

10.2.2 Increasing the Fraction of High Reputation Sellers

We now examine the impact of keeping the supply of sellers the same, and increasing the fraction of H-reputation sellers. Note that the firm can selectively incentivize H-reputation sellers to bid more through many mechanisms. It can provide monetary benefits such as lower commission rates to H-reputation sellers, and non-monetary incentives such as premium accounts and better services.

As before, we simulate all the auctions in the data a large number of times, but with the following modification – when generating the bid attributes (price, number of ratings, mean rating), for a fixed percentage of bids, instead of sampling from the full KDEs, we sample only from the top quartile of reputation values.

The results here are strikingly different from those in the previous counterfactual (see Figure 12). Here, the number of successful transactions does not increase, *i.e.*, increasing the proportion of H-reputation sellers doesn't lower the cancellation rate. While this is surprising at the first glance, there is a very simple reason for this – for the replacement bids, we sample both reputation and price from the top quartiles of the non-parametric joint distributions of bid attributes. Thus the positive correlation between reputation and prices

---

[20]Technically, in the simulation, we modify the Poisson parameters θ_p to Inf-Factor $\cdot \theta_p$.

[21]As discussed earlier, a marketplace with a healthy reputation system should see a range of prices and reputations, with L-reputation sellers receiving lower prices and H-reputation sellers receiving higher prices. Hence, the reduction in transaction price here does not indicate the failure of the reputation system; rather it is capturing how buyers' trade-off price and reputation.

[22]In our simulation, sellers don't adjust their bids to account for the increase in competition. Intuitively, if we were to allow that, transaction prices would fall even more, thereby decreasing revenues even more.

seen in the data is maintained in the simulation too.[23] Hence, from a buyer's perspective, H-reputation sellers are not necessarily 'better deals'. Thus, the number of transactions doesn't increase with more H-reputation sellers in the market. However, this is not bad news – because buyers are now replacing L-reputation, L-price sellers with H-reputation, H-price sellers in a bunch of auctions, which increases the transaction prices, and hence the revenues. Replacing 20% of random first period bids with bids from top quartile of reputation values increases revenue by 3%, and while replacing 50% of them increases the revenue by almost 6%. These are significant increases in revenue, and the fact that the number of transactions remains constant implies that these increases directly translate to profits (since transaction costs don't increase).

10.3 Modifications to Buyers' Side: Auction fees

Finally, we present a series of counterfactual experiments where we modify buyer incentives and examine the impact of these changes on buyer entry, cancellation rates, and revenues. At this point, the site does not charge the buyers any fees for posting auctions nor does it have any membership fees. To evaluate whether auction fees can improve site profits, we consider the consequences of two types of auction fees – a) a fixed fee to post an auction, and b) a percentage of MaxBid.

Introducing an auction fee has two opposing effects on revenue. On the positive side, the site has a new revenue stream. This is especially useful since it generates revenues from canceled auctions too, which were previously not at all contributing to revenues. On the flip side, auction fees now act as a negative cost on the left hand side of Equation (34), which is given by $\beta_1 \times ln(\text{auction fees} + 1)$, where β_1 is our estimate of buyers' price sensitivity. Thus, some buyers who might have entered the market and previously procured from the site, now do not even enter the auction. This leads to lower revenues from commissions for the site. We find that these two opposing forces give rise to an inverted U-shaped curve, which has a unique maxima.

Specifically, we find that fixed auction fees consistently dominate auction fees based on % of MaxBids. This is because many of the auctions in our data are low value auctions; thus fees based on % MaxBid contribute minimally to revenues. Next, we find that site revenues are maximized at an auction fee of approximately $2.75. After this point, the revenue loss from lower entry (and hence lost commissions) overwhelms the gains from auction fees. Overall, our findings suggest that the site may benefit from introducing a small fixed fee for its auctions.

11 Conclusion and Future Directions

In this paper, we develop a structural framework of buyer behavior to help researchers and managers estimate the role of seller reputations in reverse auction settings. In our framework, buyers face uncertainty over the number of bids they will receive in future and the attributes of those bids. Each period, they solve a dynamic programming problem to decide whether to terminate the auction (by choosing a submitted bid or canceling the auction) or continue to wait for another period. Unlike traditional auctions, in this setting, buyers do not pick the winning bid based on just prices; rather buyers trade-off sellers' reputations, bid prices, other bid attributes, and the cost of waiting and canceling, when making their decisions. Our framework is able to correct for dynamic selection using two types of persistent unobserved heterogeneities – in bid arrival rates and

---

[23]For all bids (both those from the full distribution and those from the top quartile of reputation), we simulate the geographic region and past interaction indicators using the logit models. So the correlations between all the bid attributes are always preserved.

in buyers' unobserved preference for bids. While persistent unobserved heterogeneity is difficult to handle in large state-space problems such as ours, we use the two-step method recently proposed by Arcidiacono and Miller (2011). In our estimation, we exploit finite dependence to reformulate value functions to improve computational tractability and then employ an EM-like algorithm to accommodate persistent unobservables.

We use our framework to study the role of reputation in online freelance marketplaces – websites that match buyers of electronically deliverable services with sellers or freelancers. While online freelancing has grown tremendously in the last few years, there exist no research studies of these marketplaces. Our estimation results from a leading online freelance place suggest that buyers are forward looking and that they place significant weight on bidder reputation. We find that not controlling for buyers' inter-temporal trade-offs and dynamic selection can considerably bias reputation estimates. Based on our estimates, we present some results on the dollar values of seller reputations and buyer entry probabilities. We also find that the site's reputation system is responsible for over 11% of its revenues. Finally, we provide three broad sets of guidelines to managers of online freelance firms. First, decreasing commission rates uniformly to either attract sellers from other sites or incentivizing all existing sellers to bid more is not a good idea. Second, it is important to incentivize high reputation sellers alone to bid more and win auctions at higher prices. Third, the introduction of a fixed $2.75 auction posting fees can increase site revenues by 5%. While our results are based on a partial equilibrium model, they nevertheless provide the best possible estimates of seller reputations and policy changes in this market.

In sum, our paper makes three clear contributions to the literature. First, from a substantive perspective, we quantify the returns to reputation in freelance marketplaces. As far as we know, this is the first paper in Marketing to study freelance marketplaces. Second, from a methodological perspective, we provide a dynamic structural framework to model and estimate the value of bidder attributes in reverse auction settings. Our framework is fairly general and can be adapted to a large class of optimal stopping problems. Third, we provide normative guidelines to managers of freelance marketplaces on improving the incentive mechanisms in their websites.

Nevertheless, there remain issues that our paper overlooks, which serve as excellent avenues for future research. First, we assume that state transitions are independent of persistent unobservables. In future, researchers may want to relax this assumption by estimating nonparametric mixture models of state transitions within the EM loop. Second, in this paper, we only look at the demand side of the freelance marketplace. While it may not be feasible to model a fully two-sided market (which would consist of a game between sellers, as well as a game between buyers and sellers), a supply side model that explores the primitives of sellers' costs and their bidding strategy would be a good next step. Doing so would allow us to run counterfactuals that take sellers' strategic behavior into consideration and further help managers optimize the commission structure on their websites. Third, we only consider buyers' behavior within an auction, and ignore inter-auction dynamics. However, we know from previous research that agents learn about market conditions over time (Crawford and Shum, 2005; Narayanan and Manchanda, 2009). Models that incorporate learning in this context would be especially useful since they would shed light on how learning affects agents' ability to build and sustain reputation in this marketplace.

Tables and Figures

Region Code	Countries
1	Indian sub-continent – India, Pakistan, Bangladesh, *etc.*
2	Developed countries – USA, Canada, Australia, Western Europe, *etc.*
3	Eastern Europe – Romania, Russia, Estonia, *etc.*
4	Everything else – Philippines, China, Korea, *etc.*

Table 1: Code for buyer and seller geographic regions.

MaxBid	Total auctions	Total bids	No. of canceled auctions	No. of auctions where buyer picked a bid	% of auctions where buyer picked a bid	No. of auctions where buyer picked lowest bid	% of uncanceled auctions where buyer picked lowest bid
10	1495	10782	348	1147	76.72 %	602	52.48 %
20	2172	18042	398	1774	81.67 %	794	44.76 %
30	1501	14030	292	1209	80.54 %	468	38.71 %
40	963	8314	232	731	75.90 %	290	39.67 %
50	2925	29169	588	2337	79.89 %	899	38.47 %
60	574	5607	144	430	74.91 %	157	36.51 %
70	245	2373	57	188	76.73 %	75	39.89 %
80	336	3632	70	266	72.67 %	95	35.71 %
90	122	1307	29	93	76.22 %	39	41.93 %
100	2999	33671	744	2255	75.19 %	853	37.83 %
Total	13332	126927	2902	10430	78.23 %	4272	40.96 %

Table 2: Summary statistics of auction outcomes

Buyer Attributes	Summary Statistics					
	Mean	Std. Dev	25^{th}	50^{th}	75^{th}	(Min, Max)
No. of ratings	40.87	125.33	2	10	34	(0, 1335)
Avg. ratings	8.3	3.57	9.54	9.96	10	(0, 10)
Avg. ratings (if rated)	9.79	0.58	9.85	10	10	(1, 10)
No. of uncanceled auctions	39.92	149.56	2	11	36	(0, 2121)
No. of canceled auctions	25.58	55.00	2	7	23	(0, 555)

Table 3: Summary statistics of buyer attributes.

Buyer Region (%)	Seller Region (%)			
	Region 1	Region 2	Region 3	Region 4
1 : 6.44	65.37	9.11	11.53	13.98
2 : 81.88	54.79	14.68	15.21	15.32
3 : 2.47	54.42	12.94	17.67	14.97
4 : 9.21	57.81	11.95	14.85	15.38

Table 4: Distributions of buyer and seller regions

	No. of periods active before choosing bid or canceling			No. of bids received		
MaxBid	Mean, Std. dev	25th, 50th, 75th percentiles	Min, Max	Mean, Std. dev	25th, 50th, 75th percentiles	Min, Max
10	2.03, 2.33	1, 1, 2	1, 14	7.30, 8.32	2, 5, 9	1, 88
20	2.12, 2.28	1, 1, 2	1, 14	8.34, 10.04	2, 5, 10	1, 140
30	2.46, 2.63	1, 1, 3	1, 14	9.32, 10.53	2, 5, 12	1, 80
40	2.52, 2.64	1, 1, 3	1, 14	8.76, 9.64	2, 6, 12	1, 74
50	2.64, 2.74	1, 1, 3	1, 14	10.02, 11.35	3, 6, 12	1, 86
60	2.89, 2.82	1, 2, 4	1, 14	9.76, 10.90	3, 6, 12	1, 88
70	3.17, 3.32	1, 2, 4	1, 14	9.59, 11.02	3, 6, 12	1, 71
80	3.40, 3.39	1, 2, 4	1, 14	10.89, 12.36	3, 6.5, 14	1, 94
90	3.54, 3.60	1, 2, 5	1, 14	10.98, 11.73	2.5, 6.5, 15	1, 59
100	3.35, 3.21	1, 2, 5	1, 14	11.19, 12.90	3, 7, 14	1, 131
Total	2.67, 2.81	1, 1, 3	1, 14	9.54, 11.09	3, 6, 12	1, 140

Table 5: Summary statistics of periods active and bids received

All Bidders						
Bidder Attributes	Summary Statistics					
	Mean	Std. Dev	25th	50th	75th	Min, Max
No. of ratings	32.32	73.65	1	6	31	0, 1679
Avg. ratings	6.94	4.14	1	9.33	9.9	0, 10
Avg. ratings (if rated)	9.23	1.3	9.03	9.66	10	1, 10
1.98% of bidders have interacted with the buyer in the past.						
Accepted Bidders						
No. of ratings	51.7	106.29	3	16	54	0, 1618
Avg. ratings	8.45	3.15	9.13	9.75	10	0, 10
Avg. ratings (if rated)	9.59	0.68	9.45	9.82	10	1, 10
10.46% of accepted bidders have interacted with buyer in the past.						

Table 6: Summary statistics of seller attributes of all bids and accepted bids.

	MaxBid									
Summary Statistics	10	20	30	40	50	60	70	80	90	100
All Bids Received										
Mean	8.77	16.44	24.24	32.07	38.71	47.25	54.45	59.26	68.46	72.52
Std. Dev	1.84	4.4	6.67	9	12.12	13.59	15.44	19.73	21.62	27.1
25th percentile	8	13	20	25	30	40	45	50	50	50
50th percentile	10	19.99	25	35	40	50	60	60	75	80
75th percentile	10	20	30	40	50	60	70	80	90	100
Min, Max	4, 10	4, 20	6, 30	4, 40	4, 50	4, 60	4, 70	4, 80	4, 90	4, 100
Accepted Bids										
Mean	8.87	16.34	24.35	31.98	37.1	47.3	56.27	60.05	67.24	65.82
Std. Dev	1.87	4.61	7.05	9.41	13.05	14.74	15.78	20.69	25.34	28.87
25th percentile	8	13	20	25	25	40	45	49	50	40
50th percentile	10	20	29	35	40	50	60	65	75	70
75th percentile	10	20	30	40	50	60	70	80	90	100
Min, Max	4, 10	4, 20	6, 30	4, 40	4, 50	4, 60	5, 70	5, 80	8, 90	4, 100

Table 7: Summary statistics of bid prices.

		Model 1		Model 2		Model 3	
Coefficients varying within the bids nest		Coefficient	Std. error	Coefficient	Std. error	Coefficient	Std. error
β_1	ln(bidprice + 1)	−0.5852	0.0255	−0.2744	0.0063	−0.6513	0.0166
β_2	Indicator no. seller ratings = 0	0.2660	0.1824	0.0738	0.0779	0.1751	0.188
β_3	ln(no. of seller ratings + 1)	0.1131	0.0094	0.0508	0.0040	0.1205	0.0095
β_4	Seller mean rating (centered)	0.0994	0.0244	0.0448	0.0107	0.1063	0.0249
β_5	Squared seller mean rating (centered)	0.0065	0.0037	0.0035	0.0017	0.0083	0.0038
β_6	Seller mean rating (centered)×	0.0709	0.0090	0.0321	0.0039	0.0762	0.0099
	ln(no. of seller ratings + 1)						
β_7	MaxBid × ln(no. of seller ratings + 1)	−0.0001	0.0001	−0.0000	0.0000	−0.0001	0.0001
β_8	Buyer mean rating×	0.0012	0.0004	0.0006	0.0002	0.0013	0.0004
	Seller mean rating (centered)						
β_9	Indicator for i and j past interaction	0.7650	0.0400	0.3337	0.0155	0.792	0.0272
β_{10}	Seller region = 1	−0.0541	0.0413	−0.0202	0.0189	−0.0480	0.0434
β_{11}	Seller region = 2	0.2792	0.0557	0.1268	0.0249	0.3010	0.0548
β_{12}	Seller region = 3	0.0926	0.0163	0.0415	0.0085	0.0986	0.0176
β_{13}	Seller region = 1 & Buyer Region = 2	0.0098	0.0580	−0.0053	0.0283	−0.0126	0.0578
β_{14}	Seller region = 2 & Buyer Region = 2	−0.0344	0.0445	−0.0202	0.0212	−0.0480	0.0455
β_{15}	Seller region = 3 & Buyer Region = 2	−0.1980	0.0522	−0.0883	0.0273	−0.2095	0.0591
β_{16}	Seller region = Buyer region $\neq$ 2	0.1390	0.0352	0.0609	0.0135	0.1446	0.0279
Coefficients common across nests							
Wait Nest							
α_{w1}	Constant			2.3629	0.1784	−0.5042	0.1811
α_{w2}	MaxBid			0.0031	0.0006	0.0024	0.0007
α_{w3}	ln(sum of buyer ratings + 1)			−0.1671	0.0154	−0.0360	0.0152
α_{w4}	Indicator no. of previous auctions = 0			−1.0086	0.1147	−0.1018	0.1140
α_{w5}	No. of uncanceled past auctions			−0.0003	0.0004	−0.0006	0.0003
α_{w6}	Cancel ratio			−1.0721	0.1004	−0.0818	0.1000
α_{w7}	Buyer tenure on site (in years)			0.0323	0.0132	0.0058	0.0132
α_{w8}	Buyer region = 1			−0.2659	0.0800	−0.0722	0.0812
α_{w9}	Buyer region = 2			0.0038	0.0620	−0.0202	0.0629
α_{w10}	Buyer region = 3			−0.0488	0.0441	−0.0415	0.0448
Bids Nest							
α_{b1}	Constant	5.2908	0.3329	3.9790	0.1232	3.0548	0.1310
α_{b2}	MaxBid	0.0025	0.0009	−0.0028	0.0007	−0.0028	0.0008
α_{b3}	ln(sum of buyer ratings + 1)	−0.3275	0.0186	−0.2449	0.0153	−0.2629	0.0161
α_{b4}	Indicator no. of previous auctions = 0	−2.7496	0.1326	−2.1635	0.1140	−3.0430	0.1212
α_{b5}	No. of of uncanceled past auctions	0.0023	0.0006	0.0013	0.0004	0.0012	0.0003
α_{b6}	Cancel ratio	−3.1816	0.1121	−2.7108	0.0958	−3.6312	0.1042
α_{b7}	Buyer tenure on site (in years)	0.0715	0.0152	0.0717	0.0014	0.1164	0.0152
α_{b8}	Buyer region = 1	−0.6953	0.1046	−0.6520	0.0954	−0.9982	0.1043
α_{b9}	Buyer region = 2	0.3869	0.0840	0.3817	0.0700	0.6210	0.0764
α_{b10}	Buyer region = 3	−0.0600	0.0523	−0.0444	0.0489	−0.0080	0.0531
α_{b11}	Indicator $s_i = 2$					3.5232	0.0436
π_1	Prob. of High type					0.2926	0.0121
δ	Discount Factor			0.2960	0.0292	0.8823	0.0314
σ	Nest correlation	0.4177	0.0169	0.1918	0.0094	0.4553	0.0126
No. of auctions = 13550							

Table 8: Estimates of Structural Parameters.

	Prob. of Cancel				
	Full data	Cohort E1	Cohort E5	Cohort E9	Cohort E13
Data	.2177	.1320	.3459	.4689	.5321
Prediction	.2177	.1979	.2363	.2786	.2932

Table 9: Model fit for the static model – the sample and model predicted probabilities of canceling are shown. Cohort Et refers to the set of buyers who **E**xited the auction in period t by choosing a bid or canceling. Note that Pr. of bid $= 1 -$ Pr. of cancel.

Cohort	Prob. of Cancel			Pr. of Bid		
(No. of obs)	Data	One type	Two types	Data	One type	Two types
S1	0.0721	0.0724	0.0532	0.4921	0.3213	0.4461
($N = 13350$)		(0.0003)	(0.0189)		(0.1708)	(0.046)
S2	0.0718	0.0744	0.0647	0.2572	0.2893	0.2949
($N = 6076$)		(0.0026)	(0.0071)		(0.0321)	(0.0377)
S3	0.0702	0.0743	0.0668	0.1854	0.2802	0.2170
($N = 4114$)		(0.004)	(0.0034)		(0.0948)	(0.0317)
S4	0.0663	0.0775	0.0778	0.1558	0.2822	0.185
($N = 3078$)		(0.0112)	(0.0113)		(0.1264)	(0.0292)
S5	0.0764	0.0779	0.0815	0.1528	0.2820	0.1576
($N = 2407$)		(0.0015)	(0.0051)		(0.1292)	(0.0048)
S6	0.0797	0.0811	0.0929	0.1444	0.2838	0.146
($N = 1857$)		(0.0014)	(0.0132)		(0.1394)	(0.0016)
S7	0.0920	0.0808	0.7921	0.1435	0.2757	0.1229
($N = 1445$)		(0.0112)	(0.0046)		(0.1322)	(0.0206)
S8	0.0869	0.0828	0.0981	0.1489	0.2813	0.1242
($N = 1105$)		(0.0041)	(0.0112)		(0.1324)	(0.0247)
S9	0.0981	0.0861	0.1083	0.1160	0.2813	0.1165
($N = 846$)		(0.012)	(0.0102)		(0.1652)	(0.0004)
S10	0.113	0.0911	0.1321	0.144	0.2945	0.1294
($N = 664$)		(0.0219)	(0.0191)		(0.1505)	(0.0145)
S11	0.1309	0.0947	0.1525	0.1474	0.2934	0.1303
($N = 489$)		(0.0362)	(0.0216)		(0.146)	(0.0172)
S12	0.1813	0.1032	0.1961	0.2249	0.3125	0.1505
($N = 353$)		(0.0781)	(0.0148)		(0.0877)	(0.0743)
S13	0.2788	0.1259	0.3621	0.2576	0.3711	0.2539
($N = 208$)		(0.153)	(0.0833)		(0.1136)	(0.0037)
S14	0.9681	0.3094	0.7026	0.032	0.6906	0.2974
($N = 99$)		(0.6587)	(0.2655)		(0.6587)	(0.2655)

Table 10: Fit of dynamic models without and with unobserved heterogeneity. Difference between model predicted probabilities and data are shown as absolute errors in brackets. Cohort St refers to the set of **S**urviving buyers at period t.
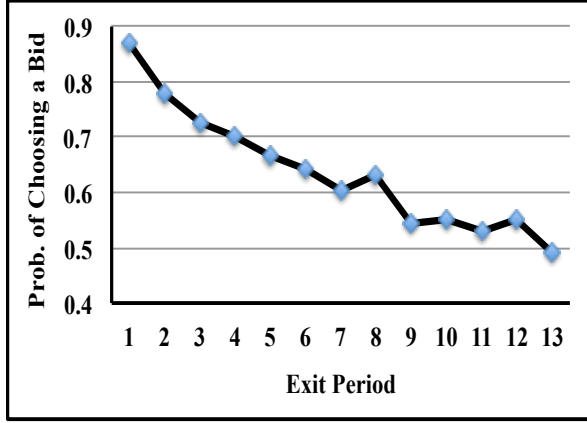
Figure 1: Probability of choosing a bid (as opposed to canceling) by buyers' exit period.
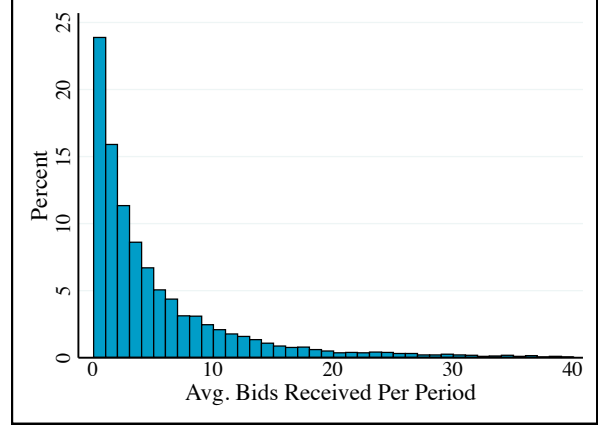

Figure 2: Histogram of average number of bids received per period by all auctions. Bin width = 1.
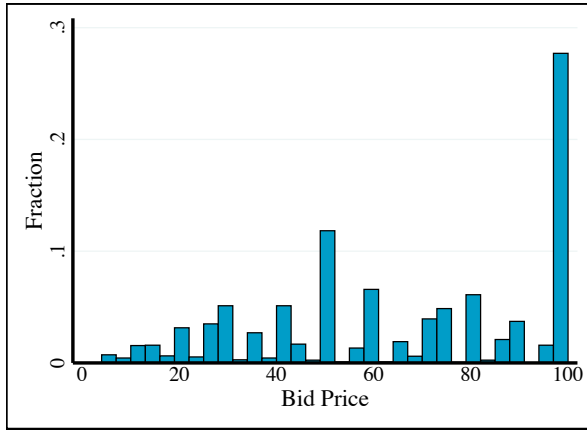

Figure 3: Histogram of prices of winning bids for MaxBid = 100. Bin width = 3.
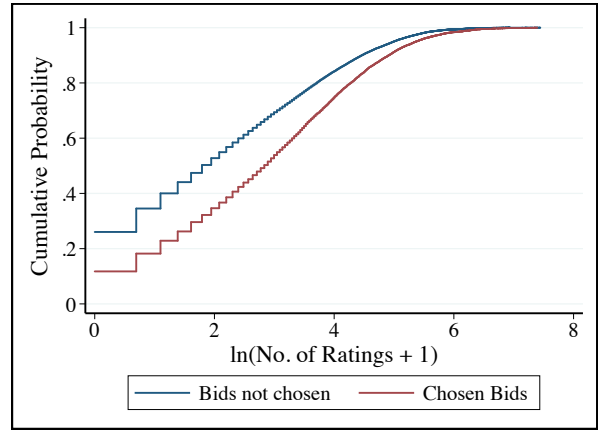

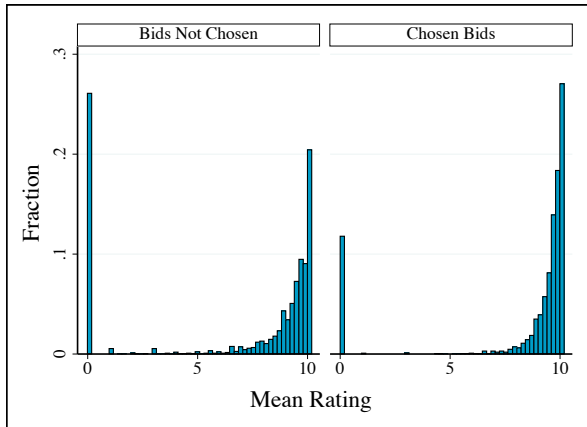Figure 4: CDFs of ln(number of ratings + 1) for accepted and rejected bidders.

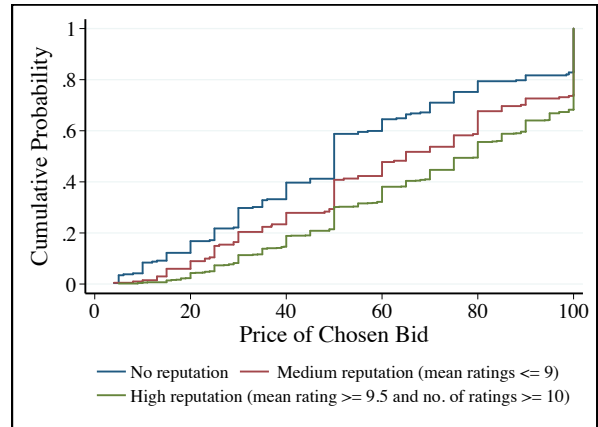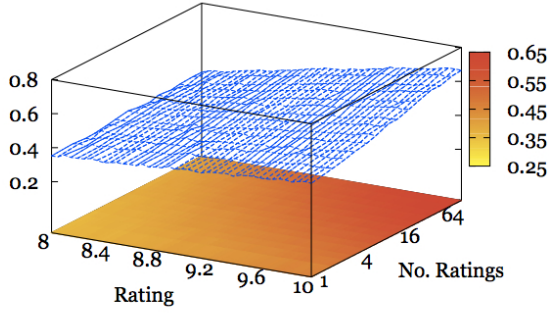Figure 5: Histogram of mean ratings for accepted and rejected bidders.
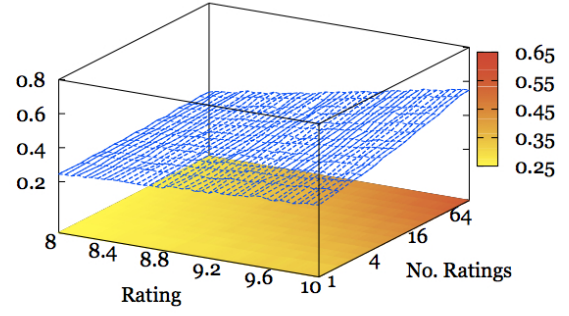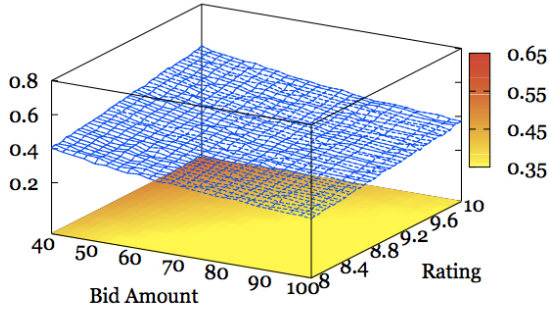

Figure 6: CDFs of prices received by sellers based on reputation. MaxBid = 100.

(a) Expected probability of winning a random auction with bid = 50.

(b) Expected probability of winning a random auction with bid = 100.

(c) Expected probability of winning a random auction when number of seller ratings is 4.
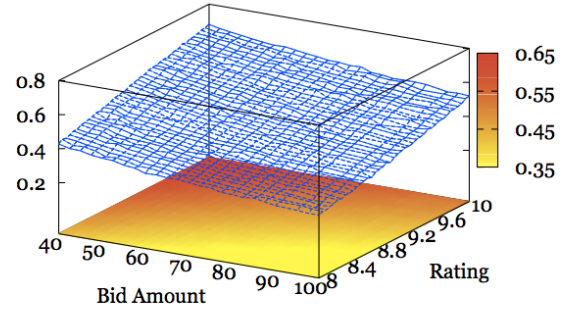
(d) Expected probability of winning a random auction when number of seller ratings is 64.
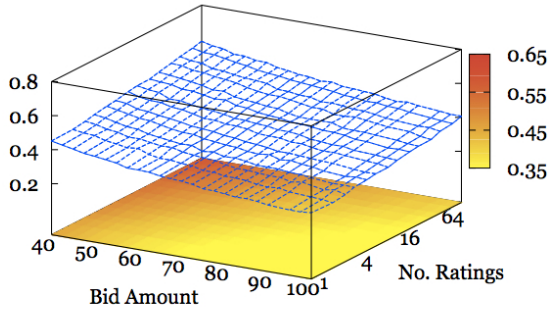
(e) Expected probability of winning a random auction when average seller rating is 9.
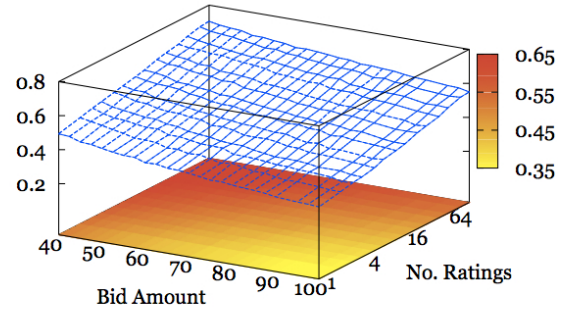
(f) Expected probability of winning a random auction when average seller rating is 10.

Figure 7: Probability of being successfully picked for a focal bidder with a certain set of reputation attributes and bid price when MaxBid = 100.
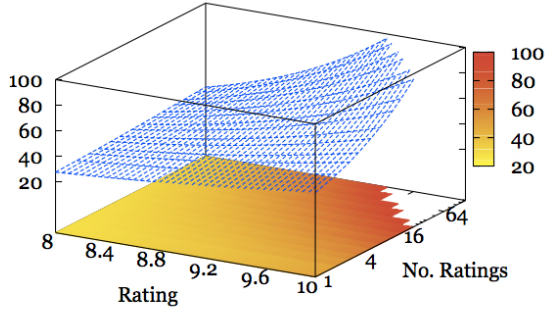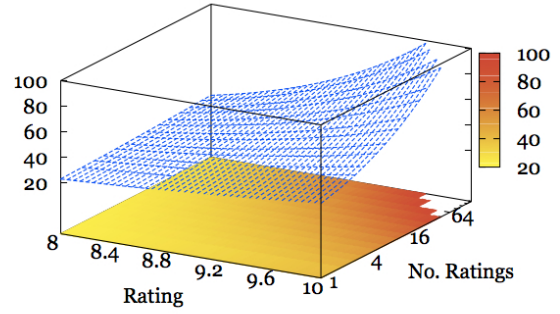
(a) Iso-class for a seller with bid = 50, avg. rating = 9.33, and no. of ratings = 6.
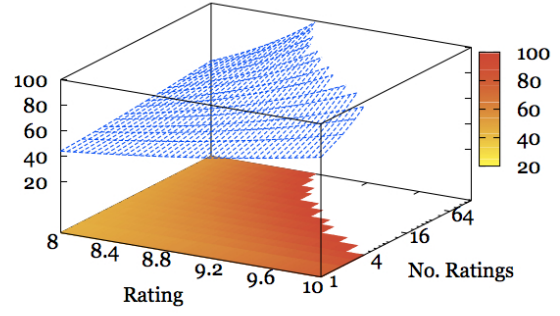
(b) Iso-class for a seller with bid = 80, avg. rating = 9.33, and no. of ratings = 6.
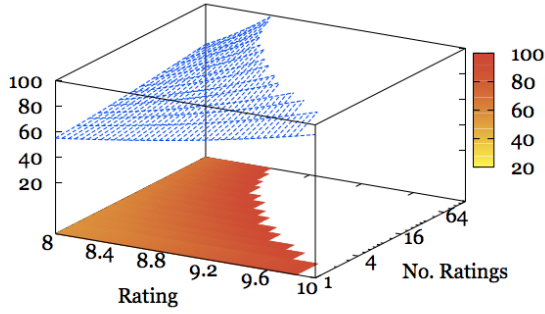
(c) Iso-class for a seller with bid = 100, avg. rating = 9.33, and no. of ratings = 6.

(d) Iso-class of sellers without past history with buyer for a focal seller with past history, bid = 25, avg. rating = 10, and no. of ratings = 31.

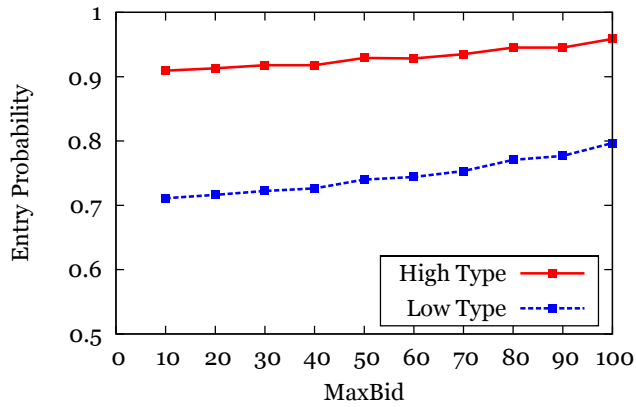Figure 8: Iso-classes for different focal sellers.

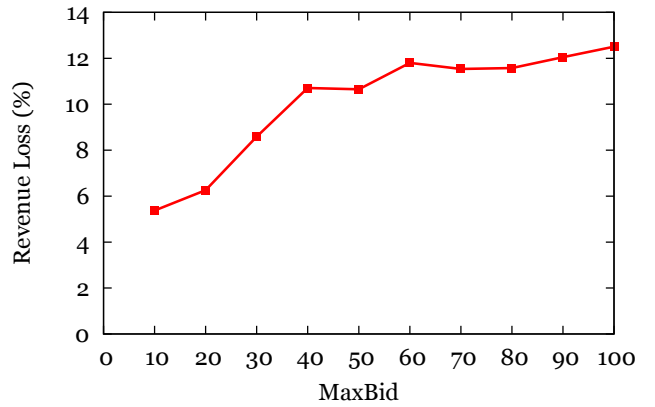Figure 9: Buyer entry probability by unobserved buyer type.

Figure 10: Percentage loss in site revenue when reputation system is switched off, by MaxBid.
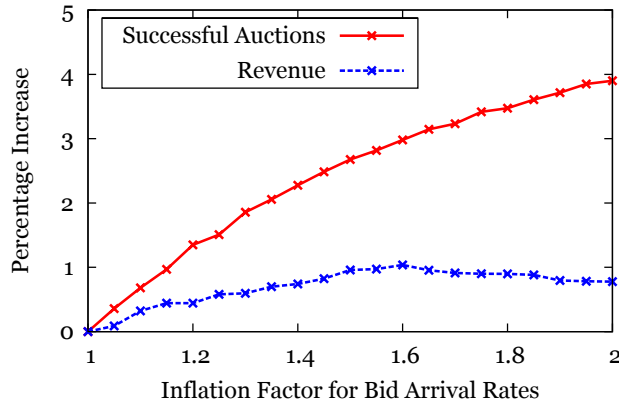
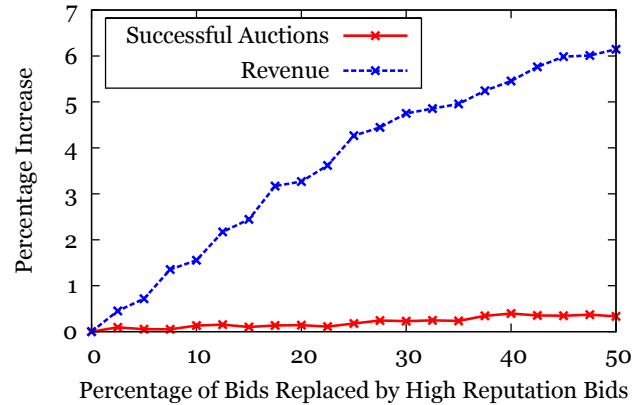Figure 11: Impact of increasing the supply of sellers.

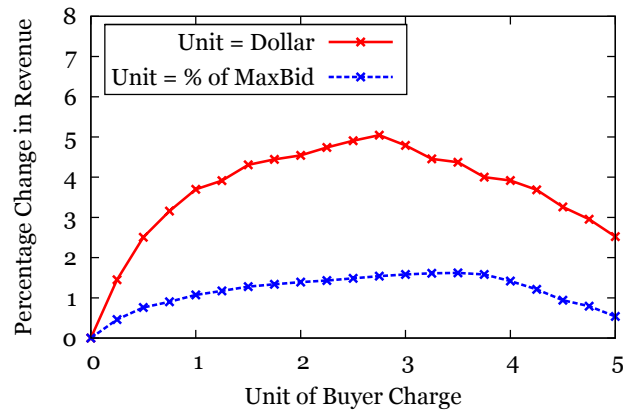Figure 12: Impact of increasing the fraction of high reputation sellers.

Figure 13: Impact of introducing auction posting fees on site revenues.

References

V. Aguirregabiria and P. Mira. Sequential Estimation of Dynamic Discrete Games. *Econometrica*, 75(1):1–53, 2007.

V. Aguirregabiria and P. Mira. Dynamic Discrete Choice Structural Models: A Survey. *Journal of Econometrics*, 156 (1):38–67, 2010.

G. A. Akerlof. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, August 1970.

G. M. Allenby and P. E. Rossi. Marketing Models of Consumer Heterogeneity. *Journal of Econometrics*, 89(1–2): 57–78, 1998.

S. Altug and R. A. Miller. The Effect of Work Experience on Female Wages and Labour Supply. *Review of Economic Studies*, 62:45–85, 1998.

P. Arcidiacono. Affirmative Action in Higher Education: How Do Admission and Financial Aid Rules Affect Future Earnings? *Econometrica*, 73(5):1477–1524, 09 2005.

P. Arcidiacono and P. B. Ellickson. Practical Methods for Estimation of Dynamic Discrete Choice Models. *Annual Review of Economics*, 3:363–394, September 2011.

P. Arcidiacono and R. A. Miller. Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity. *Econometrica*, 79(6), 2011.

P. Arcidiacono, H. Sieg, and F. Sloan. Living Rationally Under the Volcano? An Empirical Analysis of Heavy Drinking and Smoking. *International Economic Review*, 48(1), February 2007.

P. Arcidiacono, A. Khwaja, and L. Ouyang. Habit Persistence and Teen Sex: Could Increased Access to Contraception

have Unintended Consequences for Teen Pregnancies? *Journal of Business and Economic Statistics*, 30(2):312–325, April 2012.

P. Bajari and A. Hortacsu. Economic Insights from Internet Auctions. *Journal of Economic Literature*, 42(2):457–486, 2004.

P. Bajari, H. Hong, and D. Nekipelov. Game Theory and Econometrics: A Survey of Some Recent Research. Working Paper, 2010.

A. Beresteanu, P. Ellickson, and S. Misra. The Dynamics of Retail Oligopoly. Technical report, 2010.

S. Berry and P. Reiss. Empirical Models of Entry and Market Structure. *Handbook of Industrial Organization*, 3: 1845–1886, 2007.

E. T. Bradlow and Y.-H. Park. Bayesian Estimation of Bid Sequences in Internet Auctions using a Generalized Record-breaking Model. *Marketing Science*, 26(2):218–229, 2007.

E. Brousseau and J. Glachant, editors. *The Economics of Contracts: Theories and Applications*. Cambridge University Press, 2002. ISBN 9780521893138.

L. M. B. Cabral and A. Hortacsu. The Dynamics of Seller Reputation: Theory and Evidence from eBay. Technical Report 4345, C.E.P.R. Discussion Papers, 2004.

J. M. Carro and P. Mira. A Dynamic Model of Contraceptive Choice of Spanish Couples. *Journal of Applied Econometrics*, 21(7):955–980, 2006.

P. K. Chintagunta, D. C. Jain, and N. J. Vilcassim. Investigating Heterogeneity in Brand Preferences in Logit Models for Panel Data. *Journal of Marketing Research*, pages 417–428, 1991.

D. J. Chung, T. Steenburgh, and K. Sudhir. Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans. Working Paper, 2010.

G. S. Crawford and M. Shum. Uncertainty and Learning in Pharmaceutical Demand. *Econometrica*, 73(4):1137–1173, 2005.

M. W. Cripps, G. J. Mailath, and L. Samuelson. Imperfect Monitoring and Impermanent Reputations. *Econometrica*, 72(2):407–432, 03 2004.

C. Dellarocas. The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science*, 49(10):1407–1424, October 2003.

A. P. Dempster, M. N. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 39:1–22, 1977.

J. R. Douceur. The Sybil Attack. In *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, IPTPS '01, pages 251–260, London, UK, 2002. Springer-Verlag.

J.-P. Dube, G. J. Hitsch, and P. Jindal. Estimating Durable Goods Adoption Decisions From Stated Choice Data. Working Paper, 2011.

R. P. W. Duin. On the Choice of Smoothing Parameters for Parzen Estimators of Probability Density Functions. *IEEE Trans. Comput.*, 25:1175–1179, November 1976. ISSN 0018-9340.

D. H. Eaton. Valuing Information: Evidence from Guitar Auctions on eBay. Mimeo, Murray State University, 2002.

P. Ellickson, S. Misra, and H. Nair. Repositioning Dynamics and Pricing Strategy. *Journal of Marketing Research*, 2012.

P. B. Ellickson and S. Misra. Structural Workshop Paper – Estimating Discrete Games. *Marketing Science*, 30(6): 997–1010, 2011.

H. Fang and Y. Wang. Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions. 2013.

S. R. Finger. Research and Development Competition in the Chemicals Industry. Working Paper, 2008.

S. Frederick, G. Loewenstein, and T. ODonoghue. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2):351–401, 2002.

E. J. Friedman and P. Resnick. The Social Cost of Cheap Pseudonyms. *Journal of Economics & Management Strategy*, 10(2):173–199, 06 2001.

G. Gowrisankaran and M. Rysman. Dynamics of Consumer Demand for New Durable Goods. *Journal of Political Economy*, 120(8):1173–1219, 2012.

A. Gray and A. Moore. Rapid Evaluation of Multiple Density Models. In *Artificial Intelligence and Statistics*, 2003.

W. Hartmann. Intertemporal Effects of Consumption and their Implications for Demand Elasticity Estimates. *Quantitative Marketing and Economics*, 4(4):325–349, 2006.

B. Holmstrom. Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies*, 66(1):169–82, January 1999.

H. Hong and M. Shum. Using Price Distributions to Estimate Search Costs. *RAND Journal of Economics*, 37(2): 257–275, Summer 2006.

J. V. Hotz and R. A. Miller. Conditional Choice Probabilities and the Estimation of Dynamic Models. *Review of Economic Studies*, 60(3):497–529, July 1993.

A. Ihler. Kernel Density Estimation Toolbox for MATLAB, 2003. URL `http://www.ics.uci.edu/~ihler/code/kde.html`.

G. Z. Jin and A. Kato. Blind Trust Online: Experimental Evidence from Baseball Cards. Working Paper, University of Maryland, 2002.

A. Jøsang, R. Ismail, and C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. *Decision Support System*, 2007.

K. Kalyanam and S. McIntyre. Returns to Reputation in Online Auction Markets. Mimeo, Leavey School of Business and Santa Clara University, 2001.

W. A. Kamakura and G. J. Russell. A Probabilistic Choice Model for Market Segmentation and Elasticity Structure. *Journal of Marketing Research*, pages 379–390, 1989.

H. Kasahara and K. Shimotsu. Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices. *Econometrica*, 77(1):135–175, 01 2009.

J. B. Kim, P. Albuquerque, and B. J. Bronnenberg. Online Demand Under Limited Consumer Search. *Marketing Science*, 29(6):1001–1023, 2010.

S. Koulayev. Estimating Demand in Search Markets: The Case of Online Hotel Bookings. Technical report, 2009.

C. Kozierok. Online Freelancing Guide, 2011. URL http://www.freelancingguide.net/index.htm.

M. Lacity and M. Rottman. *Offshore Outsourcing of IT Work: Client and Supplier Perspectives*. Technology, Work, and Globalization. Palgrave Macmillan, 2008.

D. Lucking-Reiley, D. Bryan, N. Prasad, and D. Reeves. Pennies from eBay: the Determinants of Price in Online Auctions. *Journal of Industrial Economics*, 55(2):223–233, 06 2007.

T. Magnac and D. Thesmar. Identifying dynamic discrete choice models. *Econometrica*, 70(2):801–816, 2002.

N. G. Mankiw and P. Swagel. The Politics and Economics of Offshore Outsourcing. Working Paper 12398, National Bureau of Economic Research, July 2006.

D. McFadden. Modelling the Choice of Residential Location. In A. Karlvist, L. Lundqvist, F. Snickars, and J. Weibull, editors, *Spatial Interaction Theory and Planning Models*, volume 673, pages 75–96. North-Holland, 1978.

G. McLachlan and D. Peel. *Finite Mixture Models*. Wiley-Interscience, 2004.

M. I. Melnik and J. Alm. Does a Seller's Ecommerce Reputation Matter? Evidence from eBay Auctions. *Journal of Industrial Economics*, 50(3):337–49, September 2002.

S. Misra and H. Nair. A Structural Model of Sales-force Compensation Dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3):211–257, September 2011.

M. Morgan. The Freelance Marketplace Annual Review, 2010. WhichLance.com, 2011.

A. Murphy. A Dynamic Model of Housing Supply. Working Paper, Washington University, 2013.

S. Narayanan and P. Manchanda. Heterogeneous Learning and the Targeting of Marketing Communication for New Products. *Marketing Science*, 28(3):424–441, 2009.

W. G. Park. International Patent Protection: 1960-2005. *Research Policy*, 37(4):761–766, 2008.

Y.-H. Park and E. Bradlow. An Integrated Model for Whether, Who, When, and How Much in Internet Auctions. *Journal of Marketing Research*, November 2005.

P. Resnick and R. Zeckhauser. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In M. R. Baye, editor, *The Economics of the Internet and E-Commerce*, volume 11 of *Advances in Applied Microeconomics*. Elsevier Science, 2002.

L. Rezende. Econometrics of Auctions by Least Squares. *Journal of Applied Econometrics*, 23(7):925–948, 2008.

G. Rothwell and J. Rust. On the Optimal Lifetime of Nuclear Power Plants. *Journal of Business & Economic Statistics*, 15(2):195–208, April 1997.

J. Rust. Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, 55(5): 999–1033, September 1987.

J. Rust. Structural Estimation of Markov Decision Processes. In R. Engle and D. McFadden, editors, *Handbook of Econometrics*, volume 4, chapter 14, pages 3082–3139. Amsterdam: Elsevier-North Holland, 1994.

J. Rust and C. Phelan. How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets. *Econometrica*, 65(4):781–832, July 1997.

B. W. Silverman. *Density Estimation for Statistics and Sata Analysis*. Monographs on Statistics and Applied Probability. Chapman and Hall, 1986. ISBN 9780412246203.

M. L. Weitzman. Optimal Search for the Best Alternative. *Econometrica*, 47(3):641–54, May 1979.

R. Winkelmann. *Econometric Analysis of Count Data*. Springer, 2003. ISBN 9783540404040.

S. Yao and C. Mela. Online Auction Demand. *Marketing Science*, 27(5):861–885, 2008.

S. Yao, C. F. Mela, J. Chiang, and Y. Chen. Determining Consumers' Discount Rates With Field Studies. Working Paper, 2011.

R. Zeithammer. Forward–Looking Bidding in Online Auctions. *Journal of Marketing Research*, 43(3):462–476, 2006.

R. Zeithammer. Optimal Selling in Dynamic Auctions: Adaptation Versus Commitment. *Marketing Science*, 26(6): 859–867, 2007.

R. Zeithammer and A. Adams. The Sealed-Bid Abstraction in Online Auctions. *Marketing Science*, 29(6):964–987, 2010.

Web Appendix for "The Value of Reputation in an Online Freelance Marketplace"

1 Details of Second-Step Estimation

All the estimates from state transitions $(\theta_p, \theta_g, \theta_t)$, η_is, and non-parametric joint distributions of bid attributes are assumed to be known at this point, *i.e.*, consistent estimates of these parameters and distributions from Step 1 are used.

Throughout we estimate CCPs using weighted nested logits that are flexible functions of $\{x_{it}, \hat{\eta}_i, s_i\}$. When doing so, we treat η_is as if they were observed by substituting $\hat{\eta}_i$ in their lieu. However, because s_i is unobserved, at the each step of the EM algorithm, we weigh the nested logits with current estimates of ρs. We use flexible logits to estimate CCPs because commonly used non-parametric estimators are infeasible in our context because we have many continuous state variables and an extremely large state space. With a sufficiently large sample and an exhaustive set of covariates, CCP estimates from weighted flexible logits are similar to those obtained using non-parametric estimators (Arcidiacono and Miller, 2011; **?**). In our flexible logit, we incorporate all the state variables, their higher order terms, and interactions.

Step 1: Let the number of unobserved types be 2, and the initial guess of population probabilities be $\pi^1 = \{\pi_1^1, \pi_2^1\} = \{0.5, 0.5\}$. Let $\{\alpha^1, \beta^1, \delta^1, \sigma_1\}$ be the first guess of the structural parameters. We start with estimates from the model without persistent unobserved heterogeneity and a random α_{b11}.

Step 2: Update the ρs as follows:

$$\rho_i^2(k|d_i, x_i; \alpha^1, \beta^1, \delta^1, \sigma^1, \pi^1) = \frac{\pi_k^1 \prod\limits_{t=1}^{T_i} \left[\Pr(d_{it}|x_{it}, \hat{\eta}_i, s_i = k, \alpha^1, \beta^1, \delta^1, \pi^1)\right]^{I(d_{it})}}{\sum\limits_{k=1}^{2} \pi_k^1 \prod\limits_{t=1}^{T_i} \left[\Pr(d_{it}|x_{it}, \hat{\eta}_i, s_i = k, \alpha^1, \beta^1, \delta^1, \sigma^1 \pi^1)\right]^{I(d_{it})}} \; \forall \, k \tag{35}$$

Then update population probabilities of types as:

$$\pi_k^2 = \frac{\sum\limits_{n=1}^{N} \rho_i^2(k|d_i, x_i; \alpha^1, \beta^1, \delta^1, \sigma^1, \pi^1)}{N} \; \forall \, k \tag{36}$$

Step 3: Update CCPs using weighted flexible logits using estimates of individual weights from the previous step (ρ_i^2s) . Denote the CCP of canceling at this step as $p^2(2|x_{it}, \hat{\eta}_i, s_i = k, \rho_i^2)$ for type k.

Step 4: To set-up the weighted sum of log-likelihood from Equation (32) using choice probabilities (Equations 19-21), we need to first numerically evaluate $\int V'(x_{it+1}, \eta_i, s_i = k) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$ for each of the wait options seen in the data, for each k. We do this as follows:

- Step a: For each i, t, such that $t \neq T$, take one step into the future by generating one complete draw of a future state, say $\tilde{x}_{it+1}$, as follows:
  - Make a draw of b_{it+1} (say $\tilde{b}_{it+1}$) from the Poisson distribution $h_p(b_{it+1}|z_{t+1}, \hat{\theta}_p, \hat{\eta}_i)$
  - Make $\tilde{b}_{it+1}$ draws of the three bid attributes, bid price, number of bidder ratings, and average bidder rating, from the joint distribution $\hat{\psi}_c$, where $i \in c$. Denote each of these draws as $\{\tilde{bp}_q, \tilde{bn}_q, \tilde{br}_q\}$, where $q \in \{1, ..., \tilde{b}_{it+1}\}$.
  - For each of these draws, generate $\tilde{gx}_{iq} = \{1, I(g_i = 1), I(g_i = 2), I(g_i = 3), m_i, I(n_i = 0), \ln(n_i + 1), r_i, \ln(\tilde{bp}_q + 1), I(\tilde{bn}_q = 0), \ln(\tilde{bn}_q + 1), \tilde{br}_q\}$. Now draw a seller region for each of the q draws using the MNL distribution $h_g(\tilde{bg}_q|\tilde{gx}_{iq}, \hat{\theta}_g)$.
  - Next, let $\tilde{tx}_{iq} = \{1, I(g_i = 1), I(g_i = 2), I(g_i = 3), m_i, I(n_i = 0), \ln(n_i + 1), r_i, a_i, c_i, l_i, \ln(\tilde{bp}_q + 1), I(\tilde{bn}_k = 0), \ln(\tilde{bn}_q + 1), \tilde{br}_q, \tilde{br}_q^2, I(\tilde{g}_q = 1), I(\tilde{g}_q = 2), I(\tilde{g}_q = 3)\}$. Now draw a buyer-seller past interaction indicator for each of the $\tilde{b}_{it+1}$ draws using the logit distribution $h_t(\tilde{bt}_q|\tilde{tx}_{iq}, \hat{\theta}_t)$.
- Step b: Calculate the CCPs of canceling at $t+1$ based on the projected state $\tilde{x}_{it+1}$ as $\tilde{p}^2(2|\tilde{x}_{it+1}, \hat{\eta}_i, s_i = k, \rho_i^2)$.
- Step c: Then, use $\left[\gamma - \ln(\tilde{p}^2(2|\tilde{x}_{it+1}, \hat{\eta}_i, s_i = k, \rho_i^2))\right]$ as the current estimate of $\int V'(x_{it+1}, \eta_i, s_i = k) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$ (see Equation 23). To improve the efficiency of the results, we perform Steps a-c 100 times and take the average to obtain the current estimate of $\int V'(x_{it+1}, \eta_i, s_i = k) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$.

Step 5: Plug the current estimates of $\int V'(x_{it+1}, \eta_i, s_i = k) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$s into the choice probabilities

(Equations 19-21) and update the structural parameters by maximizing:

$$(\alpha^2, \beta^2, \delta^2) = \arg\max_{\alpha,\beta,\delta} \sum_{n=1}^{N} \prod_{t=1}^{T_i} \sum_{k=1}^{S} \rho^2(k|d_i, x_i; \alpha^1, \beta^1, \delta^1, \pi^1) \cdot \ln\left[\Pr(d_{it}|x_{it}, \hat{\eta}_i, s_i = k)^{I(d_{it})}\right] \tag{37}$$

Using the new estimates $\{\alpha^2, \beta^2, \sigma^2, \delta^2\}$, go back to Step 2 and update the ρs. Continue and repeat Steps 2-5 till all the parameters, π_ks, and the weighted log-likelihood (in Step 5) converge. At which point, we have consistent estimates of CCPs, population probabilities of unobserved types π_ks, structural parameters $\{\alpha, \beta, \sigma\}$, and the discount factor δ.

2 Identification Proof for CCPs and State Transitions

We present the proof of identification for CCPs and state transitions.

Bids arrive independently. Let $\mathcal{A}$ be the number of observable bid attributes and $\mathcal{B}$ the maximum number of bids that an auction can receive. Let $\mathcal{T}$ be the maximum time periods for which an auction can remain active, and let $\mathcal{X}$ and $\mathcal{S}$ be the total number of observed and unobserved buyer/auction specific state variables (which remain constant during the auction). Given the structure of our problem (*i.e.*, waiting is the only continuation decision), the number of possible sequences of decision and state variables is:

$$\text{Seq} = \mathcal{X} \sum_{b=0}^{\mathcal{B}} \frac{(b + \mathcal{T} - 1)!}{b!(\mathcal{T} - 1)!} \mathcal{A}^b - 1 \tag{38}$$

where $\frac{(b+\mathcal{T}-1)!}{b!(\mathcal{T}-1)!}$ is the number of ways to distribute b bids among $\mathcal{T}$ periods. Without placing any structure, the number of first step parameters to be estimated include the following number of CCPs, bid arrivals[24], joint distributions of bid attributes, and the initial distributions of unobserved types:

$$\text{CCP Par} = \mathcal{X}\mathcal{S}\mathcal{T} \sum_{b=0}^{\mathcal{B}} (b+1)\mathcal{A}^b$$
$$\text{Initial dist. of types Par} = \mathcal{X}\mathcal{S} - 1$$
$$\text{Bid arrival Par} = \mathcal{X}\mathcal{S}\mathcal{T}$$
$$\text{Joint bid dist. Par} = \mathcal{X}(\mathcal{A} - 1)$$

Hence the total number of parameters is:

$$\text{Total Par} = \mathcal{X}\mathcal{S}\mathcal{T} \sum_{b=0}^{\mathcal{B}} (b+1)\mathcal{A}^b + \mathcal{X}\mathcal{S} + \mathcal{X}\mathcal{S}\mathcal{T} + \mathcal{X}(\mathcal{A} - 1) - 1 \tag{39}$$

For identification, we need:

$$\sum_{b=0}^{\mathcal{B}} \frac{(b + \mathcal{T} - 1)!}{b!(\mathcal{T} - 1)!} \mathcal{A}^b \geq \mathcal{S}\mathcal{T} \sum_{b=0}^{\mathcal{B}} (b+1)\mathcal{A}^b + \mathcal{S} + \mathcal{S}\mathcal{T} + \mathcal{A} - 1$$
$$\Rightarrow \sum_{b=0}^{\mathcal{B}} \left[\frac{(b + \mathcal{T} - 1)!}{b!(\mathcal{T} - 1)!} - \mathcal{S}\mathcal{T}(b+1) \right] \mathcal{A}^b \geq \mathcal{S} + \mathcal{S}\mathcal{T} + \mathcal{A} - 1 \tag{40}$$

The left hand side terms have $\mathcal{A}^b$, where $\mathcal{A}$ is very large and $\mathcal{A}^b$ increases exponentially with b. Moreover, the terms within the bracket is positive for $\mathcal{T} \geq 3$ and sufficiently large bs. So as long as we have enough time periods and a large number of potential bids (B), we have more decision-state sequences than parameters. In our case, we have $\mathcal{T} = 14$ and $B = 149$, and even the mean number of bids received by buyers is around 10, so the CCPs and state transitions are non-parametrically identified.

Before we proceed though, we do need to ensure that a key corner case is identified. The corner case in these kinds of identification proofs is always the one where both time and state space dimension are lowest. In our case, since we use CCPs to calculate continuation values of waiting from period 1, we need all CCPs from second period

---

[24]If bid arrivals were not independent, the number of parameters to be estimated for the bid arrival process would be $\mathcal{X}\mathcal{S}\mathcal{T}(\mathcal{B} - 1)$ instead of $\mathcal{X}\mathcal{S}\mathcal{T}$.

forward to be identified. Hence, the corner case we check is $t = 2$ and one bid. When a buyer has one bid in the second period, there are only two ways for him to have reached that position – a) receive no bids in the first period and wait, and receive one bid in the second period, and b) receive one bid in the first period and wait, and receive no bids in the second period. Thus to calculate terminal CCPs for the second period in this state (one bid), we only have two equations. With two equations, we can identify CCPs for at most two unobserved types, but no more. Hence, this corner case establishes that, in this setting, we can only allow for two unobserved types.

Finally, note that even though the number of sequences is much greater than parameters, first period termination CCPs (CCPs for cancellation and bid choice) are not identified because there is only one path to reach a state in the first period and once a buyer terminates the auction, there are no paths out of it. Hence, these CCPs don't appear in more than one sequence and cannot be identified. However, we don't use first period CCPs in our analysis; only CCPs from the second period and forward are used in continuation value calculations. Hence the non-identification of first period termination CCPs is not a problem.[25]

2.1 Simulation Details: Aggregate Choice Probability Maps

Consider a focal bid j with fixed reputation attributes and price, say $\{bp_j, bn_j, br_j\}$. Simulate each auction i 100 times using the steps described below.

1. For each iteration of auction i, start with $t = 1$ and simulate each period as follows:
   - Draw the number of bid arrivals for this period using the Poisson estimates and $\hat{\eta}_i$. If the first bid for this iteration of the auction arrives in this period, treat it as the focal bid with reputation and price $\{bp_j, bn_j, br_j\}$. Then, draw geography and past interaction indicators for it using the logit models.
   - Then draw bid attributes for each non-focal bid that arrives in this period using the estimated KDEs and logit models.
   - If $t \neq T$, calculate the continuation value and derive the probability of picking the focal bid (if it exists), picking one of the other bids (if they exist), waiting, or canceling the auction using the closed form expressions for each unobserved type $s_i = \{1, 2\}$. Then weight the probabilities using ex post probabilities of belonging to the two unobserved types (*i.e.*, ρ_is from the last iteration of the EM algorithm).
   - Construct a CDF using the above probabilities, and draw a random number from 0 to 1 to simulate the outcome for this period.

2. Record the outcome for this period as follows:
   - If focal bid was chosen, record this, and terminate this iteration of auction i.
   - If another bid was chosen or if auction was canceled, record this, and terminate this iteration of auction i.
   Then, start next iteration for auction i if 100 iterations haven't been reached, else move on to the next auction.
   - If wait is chosen, proceed to the next time period for this iteration of auction i.

3. When 100 iterations for auction i are completed, calculate the probability of the focal bid winning auction i as = (number of times it was picked in the 100 iterations)/100.

4. When all the auctions have been simulated and the average probability of being picked in each auction has been calculated, sum the probabilities, and divide by the total number of auctions to get the average success probability for the focal bid in the website.

3 Details of Backward Solution Method used in Counterfactuals

All the estimates from state transitions $(\theta_p, \theta_g, \theta_t)$, η_is, non-parametric joint distributions of bid attributes, and the structural parameters $(\alpha, \beta, \delta, \sigma, \pi)$ are treated as known at this point.

First, split the observed state variables at time t, x_{it}, into auction/buyer specific variables that time invariant for the duration of the auction, xa_i, and bidder specific state variables that are denoted as xb_{it}. Then based on the definition in the paper, write out the inclusive value of the bid nest for buyer i at time t as: $I_{it} = I(xb_{it}) = \ln \left[\sum_{q=1}^{B_{it}} e^{\frac{Y_{iq}(xb_{it})}{\sigma}} \right]$, where $Y_{ij}(xb_{it})$ is defined in Equation (6) and B_{it} is the total number of bids that i has at period t. Next, discretize the state variables, xa_i, I_{it}, and η_i by allowing them to vary from zero to a high enough value, in steps of Δ, such that: $xa_i \in \{0, xa_{max}\}$, $I_{it} \in \{0, I_{max}\}$, and $\eta_i \in \{0, \eta_{max}\}$. Using these definitions, specify the continuation value as a function of the time-invariant observables and the inclusive value at time t as $V'(t, ax_i, I_{it}, \eta_i, s_i)$. Now to generate the value functions at each combination, follow the steps given below:

Step 1: Start at the last time period T. Since waiting is not an option at T, there is no continuation value in this period.

---

[25]Unlike first period termination CCPs, first period wait CCPs are identified because they appear in many decision-state paths.

Hence, the conditional choice probability of cancellation at T is simply obtained from the following equation:

$$P(2|T, xa_i, I_{iT}, \eta_i, s_i) = \frac{1}{1 + e^{[W_{ib}(xa_i, s_i) + \sigma I_{iT}]}} \tag{41}$$

Now using all the values in the grid of a_i and I_{iT} (in $\{0, xa_{max}\}$ and $\{0, I_{max}\}$ in steps of Δ), populate the cancellation probability matrix. Note that at this point we can simply ignore η_i and s_i since we are just calculating the cancellation probabilities at different levels of a_i and I_{iT} at T.

Step 2: Next, use the matrix of cancellation probabilities at T to calculate the matrix of continuation values at each point in the grid of $\{xa_i, I_{iT}\}$ using the following equation:

$$V'(T, xa_i, I_{iT}, \eta_i, s_i) = \gamma - \ln(P(2|T, xa_i, I_{iT}, \eta_i, s_i)) \tag{42}$$

Step 3: Now move to period $t = T - 1$ and at each grid point in $\{a_i, I_{iT}, \eta_i\}$, take the following steps:

- Step 3a: Follow Step a under Step 4 in §1 of the Technical Appendix in order to simulate one future draw of bids for period $t + 1$. Based on the draw of future bids, obtain the new inclusive value I_{it+1} (using both current bids and simulated bids). Then look up the matrix of continuation values calculated in the previous step and obtain $V'(t + 1, xa_i, I_{it+1}, \eta_i, s_i)$
- Step 3b: Repeat Step 3a for a large number of times to obtain $\int V'(t+1, xa_i, I_{it+1}, \eta_i, s_i) f(x_{it+1}|d_{it}, x_{it}, \eta_i) dx_{it+1}$.
- Step 3c: Now plug back this integral into the choice-specific continuation value of waiting at t to obtain:

$$v(1, t, xa_i, I_{it}, \eta_i, s_i) = W_{iw}(xa_i) + \delta \int V'(t + 1, xa_i, I_{it+1}, \eta_i, s_i) f(x_{it+1}|1, x_{it}, \eta_i) dx_{it+1}$$

- Step 3d: Substitute this choice specific value function and the parameter estimates in Equation 20 to obtain the cancellation choice probability as:

$$P(2|t, xa_i, I_{it}, \eta_i, s_i) = \frac{1}{1 + e^{[v(1,t,xa_i,I_{it},\eta_i,s_i)]} + e^{[W_{ib}(xa_i, s_i) + \sigma I_{it}]}} \tag{43}$$

Performing Steps 3a-3d at all the grid points in $\{xa_i, I_{it}, \eta_i\}$ should give a matrix of cancellation probabilities at t. Then, as in Step 2, use the matrix of cancellation probabilities at t to calculate the matrix of continuation values at each point in the grid of $\{xa_i, I_{it}, \eta_i\}$ using the following equation:

$$V'(t, xa_i, I_{it}, \eta_i, s_i) = \gamma - \ln(P(2|t, xa_i, I_{it}, \eta_i, s_i)) \tag{44}$$

Step 4: Repeat Step 3 for all periods till the algorithm reaches $t = 1$ and obtain the continuation values $V'(t, xa_i, I_{it}, \eta_i, s_i)$ $\forall\, 1 \leq t < T$.

Step 5: With the matrices of continuation values derived above and the estimated parameters, the system can now be simulated as many times as necessary to generate the counterfactual outcomes.

4

4 Additional Tables and Figures

Project Title	MaxBid	Project Description
Web Site Photo Gallery Assistance	20	Have code already for a website's photo gallery. And for the website, I just need gallery code merged into existing page.
Need Recommendations for AdSense placement in our website	50	Hi, We have placed AdSense Ads in our website, even though we are getting some revenue from AdSense, we are not getting return on our investment. So, we would like to try placing AdSense Ads in strategic places and need recommendations. We need recommendations with size of Ads, location of Ad placement and any other help that you can give us to increase our revenue through Ads. Thanks
Wordpress Theme Modification	80	Hi Guys, I got the "TheCorporation" theme from elegant designs. I have a PSD ready on how the theme needs to look like. I also have the original PSD files of the theme. I need you to edit the theme so that it will look exactly the same as in the PSD that it needs to look like.
Streaming Data Capture	100	I am looking for somone [sic] who can write me a program that can store streaming numerical data on a web page to an excel spreadsheet, or as csv or at least as a text file.

Table 11: Examples of project titles and descriptions

z_t	θ_p	std. error
Ind. $t = 1$	4.4866	0.0445
Ind. $t = 2$	2.6337	0.0447
Ind. $t = 3$	1.8599	0.0455
Ind. $t = 4$	1.3820	0.0468
Ind. $t = 5$	1.0787	0.0484
Ind. $t = 6$	0.8096	0.0500
Ind. $t = 7$	0.6770	0.0512
Ind. $t = 8$	0.4983	0.0518
Ind. $t = 9$	0.3739	0.0525
Ind. $t = 10$	0.3037	0.0550
Ind. $t = 11$	0.2351	0.0552
Ind. $t = 12$	0.1358	0.0571
Ind. $t = 13$	0.0513	0.0572
No. of auctions, auctions×timeperiods = 6076, 28817		
Goodness of Fit Measures RMSE, MAE, ME = 0.9258, 0.4249, 3.59×10^{-8}		

Table 12: Estimates of θ_p from the fixed effects Poisson model. Standard errors are adjusted for clustering on auction identity. All the estimates, except the last two, are significant at the 1% level (the parameter coefficient for Ind. $t = 12$ is significant at the 10% level).

gx_{ij}	Region 1		Region 2		Region 3	
	θ_{g1}	Std. error	θ_{g2}	Std. error	θ_{g3}	Std. error
Buyer region = 1	0.2094	0.0426	−0.1800	0.0597	−0.1322	0.0561
Buyer region = 2	−0.0496	0.0288	0.2007	0.0384	−0.0042	0.0365
Buyer region = 3	0.0051	0.0185	0.0382	0.0242	0.0563	0.0226
MaxBid	−0.0006	0.0004	−0.0000	0.0005	0.0025	0.0005
Ind. no. of buyer ratings = 0	−0.3433	0.1531	0.7286	0.2137	−0.3328	0.1890
ln(no. of buyer ratings + 1)	0.0236	0.0068	0.0324	0.0087	−0.0260	0.0086
Buyer mean ratings	−0.0338	0.0154	0.0660	0.0215	−0.0276	0.0190
ln(bidprice + 1)	0.0654	0.0184	0.1120	0.0236	−0.0861	0.0229
Ind. no. of seller ratings = 0	−1.6200	0.0809	1.1569	0.1232	3.7824	0.1554
ln(no. of seller ratings + 1)	−0.0046	0.0067	−0.1334	0.0087	0.0376	0.0082
Seller mean rating	−0.1732	0.0085	0.1491	0.0098	0.4270	0.0159
Constant	3.0200	0.1789	−2.4263	0.2535	−3.5782	0.2504
No. of observations = 126927						
Log likelihood = −145597.83						

Table 13: Estimates of $\theta_{g1}, \theta_{g2}, \theta_{g3}$ from the Multinomial Logit model. Region 4 is the base outcome.

tx_{ij}	θ_t	Std. error
Buyer region = 1	−0.0041	0.1277
Buyer region = 2	0.0048	0.0813
Buyer region = 3	−0.0363	0.0616
MaxBid	−0.0109	0.0014
No. of uncanceled past auctions	0.0007	0.0001
Cancel ratio	−0.8858	0.1050
Buyer tenure on site (in days)	−0.0006	0.0000
Ind. no. of buyer ratings = 0	−1.2024	0.4578
ln(no. of buyer ratings + 1)	0.5432	0.0220
Buyer mean ratings	−0.0165	0.0513
Seller region = 1	0.1341	0.0638
Seller region = 2	0.3069	0.0777
Seller region = 3	0.0115	0.0258
ln(bidprice + 1)	0.2271	0.0589
I.(no. of seller ratings = 0) = 1	0.1668	0.9083
ln(no. of seller ratings + 1)	0.4639	0.0153
Seller mean rating	−0.1491	0.2145
Square of seller mean rating	0.0323	0.0129
Constant	−7.4925	1.0504
No. of observations = 126927		
Log likelihood = −9671.063		

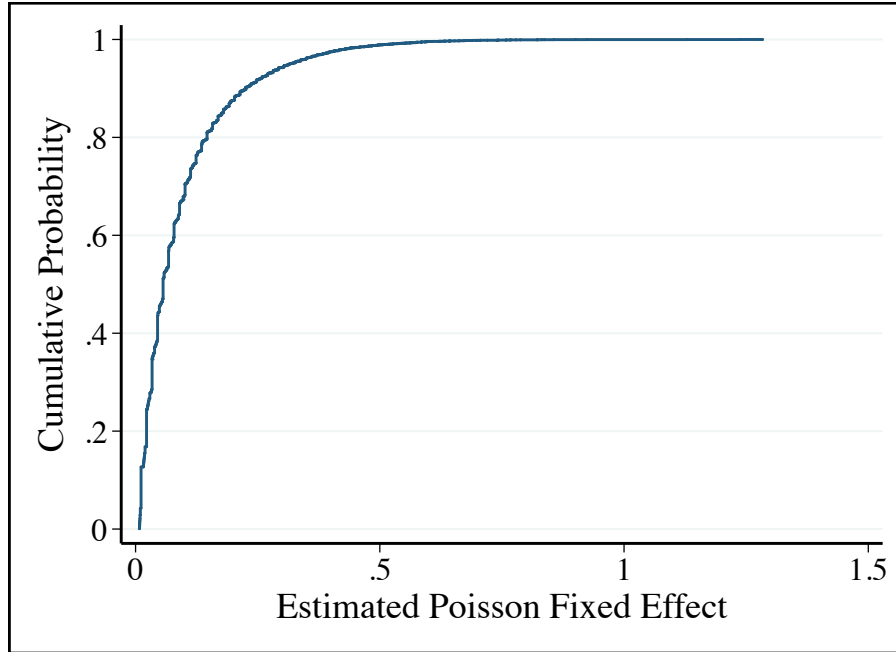Table 14: Estimate of θ_t from the Logit model. $bt_{ij} = 0$ is the base outcome.

Figure 14: CDF of estimated Poisson fixed effects ($\hat{\eta}_i$s) for the auctions in our sample. Summary statistics of the distribution: Mean, Std. Dev, 25[th], 50[th], 75[th] percentiles = 0.0952, 0.1068, 0.0246, 0.0563, 0.1238.