

HOW EFFICIENT ARE DECENTRALIZED AUCTION PLATFORMS?

AARON L. BODOH-CREED - UNIVERSITY OF CALIFORNIA, BERKELEY

JÖRN BOEHNKE - HARVARD UNIVERSITY

BRENT HICKMAN - UNIVERSITY OF CHICAGO

ABSTRACT. We provide a model of a decentralized, dynamic auction market platform (e.g., eBay) in which a continuum of buyers and sellers participate in simultaneous, single-unit auctions each period. Our model accounts for the endogenous entry of agents and the impact of intertemporal optimization on bids. We estimate the structural primitives of our model using Kindle sales on eBay. We find that just over one third of Kindle auctions on eBay result in an inefficient allocation with deadweight loss amounting to 14% of total possible market surplus. We also find that partial centralization - for example, running half as many 2-unit, uniform price auctions each day - would eliminate a large fraction of the inefficiency, but yield slightly lower seller revenues. Our results also highlight the importance of understanding platform composition effects - selection of agents into the market - in assessing the implications of market design. We close by proving that the equilibrium of our model with a continuum of buyers and sellers is an approximate equilibrium of the analogous model with a finite number of agents.

Keywords: Dynamic Auctions, Approximate Equilibrium, Internet Markets.

JEL Codes: C73, D4, L1

Date: January 2017.

We thank (in no particular order) Dan Quint, Michael Dinerstein, Lanier Benkard, Alex Torgovitsky, and John van den Berghe for helpful comments and feedback. We also thank seminar participants at the University of Wisconsin–Madison and University of Chicago, as well as participants of the 2014 IIOC Conference (Northwestern University), the 2014 Conference on Auctions, Competition, Regulation, and Public Policy (Lancaster University), 2015 Econometric Society World Congress (Montreal), 2015 INFORMS Conference (Philadelphia), and the 2016 Econometric Society North American Summer Meetings (University of Pennsylvania). We are also greatly indebted to Steven Mohr and Vincent Johnson (UChicago SSCS), without whose dedicated technical support this project—from web scraping to estimation—would not have been possible. Jörn Boehnke gratefully acknowledges the support of the Center of Mathematical Sciences and Applications at Harvard University during the course of this work.

1. INTRODUCTION

Online market platforms are increasingly important in today's economy, and the goal of these platforms is to provide a venue for buyers and sellers of various goods to transact. For example, in 2014 eBay reported USD\$82.95 billion in sales volume and 8.5% annual growth after nearly two decades in business.¹ StubHub, a platform for selling tickets to events such as soccer games, and Upwork, a platform for recruiting freelance workers, each host annual transaction volumes in the hundreds of millions or billions of dollars. Since a large number of participants are exchanging a broad array of products on these platforms, each platform has sophisticated search tools to help buyers and sellers find partners to transact with.

Given the power of modern search algorithms and the thickness of the markets, one might conjecture that these platforms would do an excellent job of matching buyers and sellers, eliminating market frictions, and generating efficient trade. This conjecture is particularly compelling in cases where the products on offer are homogenous and buyer and seller reputation are not significant barriers to trade. Our goal is to test this conjecture by estimating a novel model of the eBay auction platform using data on sales of new Amazon Kindle Fire tablets.

On the eBay platform a large number of participants compete in a large number of auctions each day, and buyers and sellers can participate across successive days. In this paper we provide a rich model of such an auction platform in which a continuum of buyers is matched to a continuum of seller auctions each period. After matching has occurred, each single-unit auction is executed independently, auction winners (and the associated sellers) exit the market, losing bidders move on to the next period, and new bidders enter at the end of each period. We include a costly endogenous entry decision to capture the time and effort costs of participation. We use an extensive dataset on new Amazon Kindle Fire tablets to estimate the structural model primitives such as the matching process that allocates potential buyers to auction listings, the monetized cost of participation, and the steady-state distributions of buyer valuations and seller reserve prices. While the participation cost we find is low, on the order of \$0.10, it turns out to be an important regulator of the number and types of buyers in the market.

Having estimated the structural model, we move on to investigate the efficiency of the market. Platform markets like eBay exist for the purpose of reducing frictions that impede trade, so it is natural to assess how closely eBay approaches the ideal of fully efficient trade. To fix ideas, suppose there are exactly two listings and four bidders, two bidders with high valuations and two bidders with low valuations. The social planner's

¹Information downloaded from <https://investors.ebayinc.com/secfiling.cfm?filingID=1065088-15-54&CIK=1065088> on 11/17/2015.

preferred outcome is one where each auction attracts one high value bidder as this would guarantee that the high value bidders win in any monotone bidding equilibrium. However, when there is randomness in the bidder-listing match process, there will be a positive probability that one auction listing will not have a high-value participant, meaning a low value bidder wins at a low price and a high-value bidder loses. Another way of putting it is that the matching frictions mean some auctions have too much competition and others too little relative to an efficient allocation.

We begin our counterfactual analysis by using two separate methods to measure inefficiency under the current market conditions. Our first method relies heavily on the raw data. Using our estimated buyer-seller ratio we can count the number of times in our data that a bidder with an inefficiently low value (i.e., low bid) won an auction and prevented a high value bidder from receiving that item. This method can only give a lower bound on the prevalence of inefficient allocations because, for example, it cannot detect scenarios where multiple high-value losers attended the same auction. We find that within the listings for new Kindles, at least 27.6% of all auctions allocate goods to buyers with inefficiently low valuations.

In our second method, we use the structural estimates to get a precise value for the fraction of auction listings that award an object to a buyer whose private value is inefficiently low. This method also allows us to quantify the deadweight loss, which is defined as the average difference in value between high-value losers and low-value winners. We calculate that 36% of used Kindle listings result in inefficient allocations and that the total deadweight loss amounts to roughly 14% of potential market surplus. In other words, we estimate that eBay is able to achieve 86% of all possible gains from trade.²

Next, we explore the implications of alternative spot-market mechanisms eBay could use to improve efficiency. Specifically, we consider the welfare cost of eBay's choice to use single-unit auctions, which we refer to as *decentralization*. We use our estimates to analyze outcomes of alternative markets where, instead of single-unit auctions, eBay runs K -unit, uniform price auctions, and we use K as a measure of the market's *centralization*. The most extreme version of this counterfactual would be a single multi-unit, uniform-price auction each day. We find that aggregating auctions together so that eBay runs half as many auction listings for 2-units each day recovers 35% of the welfare loss by improving the efficiency of the allocation, while running a quarter as many 4-unit auctions recovers over 57% of the welfare loss. However, centralization reduces the expected sale

²In the main text we also show that using a completely random allocation achieves 47% of the maximum possible welfare. The status quo outcome achieves only 74% of the gains of trade relative to the random assignment benchmark.

price, which in turn reduces eBay's revenues.³ In addition to being a vehicle for analyzing the welfare losses, we believe that centralizing auctions is a practical design strategy in settings wherein the goods are homogenous (e.g., new Kindles).

While we conduct our estimates within the particular eBay context, we believe that the degree of welfare losses we find should temper optimistic expectations of the beneficial effects of platform markets more generally. In addition, the market centralization solution we propose is far more broadly applicable than a single, isolated eBay market, and we believe it could be worthwhile considering similar centralization-oriented designs in other platform market contexts. For example, centralization may be possible for standardized back-office tasks that are bought and sold on Upwork.

For our second counterfactual exercise, we use our estimated model to tease apart the welfare effects of a change in the market design into a component capturing changes in the dynamic incentives (i.e., changes in the continuation value) and a second component reflecting changes in the platform composition (i.e., the ratio of buyers to sellers and the distribution of buyer types). Our empirical approach is necessary because these components interact in complex ways in our model. For example, any change in the market structure that increases the continuation values of the buyers (e.g., making the allocation more efficient) will increase bid shading. At the same time, increased efficiency may encourage low value buyers to leave the market as the probability of a low value buyer winning an auction drops. As low value buyers exit, the typical bidder has a higher value for the good, meaning that competition from the typical opponent bidder may become more intense. These countervailing effects make it unclear whether a platform must necessarily benefit from improvements in market efficiency.

In our counterfactuals we consider an increase of the participation cost for bidders. An increase in participation costs pushes low value bidders out of the market (i.e., a platform composition effect), altering the steady state distribution of bidder values and the ratio of buyers to sellers. The cost increase also has an effect on dynamic incentives by reducing the continuation values (and hence raising the bids) of the buyers. We find that the platform composition effects have from two to ten times more effect on market efficiency than the dynamic incentive effects. Although we consider a particular change in the market structure, our analysis emphasizes the importance of attending to the selection of users into the platform when redesigning any platform market.

For our final counterfactual exercise, we use our model to investigate optimal seller reserve prices. For a seller who values the object at \$0, the optimal reserve price for the seller ought to be no less than the expected revenue from selling the item the following

³Currently eBay charges about 10% of the revenue generated by the Kindle auctions.

day multiplied by the seller's discount factor. Given that more than two thirds of observed reserve prices are close to zero, most sellers are not setting the reserve price to account for the opportunity cost of selling the good today, which creates a puzzle. However, given the intensity of bidder competition in the market, we show that the benefit to the seller of setting an optimal dynamic reserve price is on the order of \$1, which suggests that sellers have weak incentives to choose the optimal reserve price.⁴

On the empirical front, we make several contributions to the literature on identifying auction models. A key feature of our estimator is that it requires only observables that are readily available on many platform websites. In particular, we are able to identify the average number of buyers per auction without assuming that we observe all of the bidders in each auction. If bid submission times are randomly ordered, then some auction participants with an intent to bid may be prematurely priced out of the spot market before they get a chance to submit their bid. Therefore, the total number of unique bidders within a given eBay auction constitutes a lower bound on the actual number of competitors. Our nonparametric identification argument for the dynamic structural model requires only that we observe this lower bound on the number of competitors, the seller reserve price, and the highest losing bid within each auction.

Our identification strategy also lets us separately identify bid shading (i.e., bidding strictly below one's private valuation) due to the use of a nontruthful pricing rule (e.g., a first-price auction) and bid shading due to intertemporal incentives. From a buyer's perspective, failure to win an auction today is no tragedy since he can return tomorrow and bid again, which implies there is an opportunity cost to winning today. The opportunity cost determines the degree of demand shading within the current period. Given a value for the time discount factor, we show that this demand shading factor is nonparametrically identified from observables that are readily available on eBay. We also show that when the spot-market pricing rule is non-second-price—so that the winner's bid may directly affect the sale price—then the additional, static demand shading incentive is layered on top of the dynamic demand shading incentive in an intuitive way that allows for straightforward econometric identification. This is important since many electronic auction pricing rules (including eBay's) are known to deviate from the standard second-price form in empirically relevant ways.

Our empirical analysis reveals that the degree of bid shading incentivized through intertemporal opportunity costs is significantly larger than the demand shading generated by the choice of a nontruthful spot market auction mechanism. In other words, it is more

⁴A similar result was found using very different techniques by Einav, Kuchler, Levin, and Sundaresan [2015].




FIGURE 1. Static Versus Dynamic Demand Shading Incentives

important for bidders to understand intertemporal opportunity costs than how to strategically bid under a nontruthful pricing rule. To make this concrete, Figure 1 plots the equilibrium bidding strategies under first and second price auction rules in a dynamic auction market given the economic primitives we estimate. The 45 degree line can be interpreted as the equilibrium of a static second price auction that omits the dynamic opportunity cost. The bid shading caused by the intertemporal opportunity cost is represented by the difference between the bidding strategy under the second price auction (SPA) mechanism in a dynamic setting and the 45 degree line. The additional demand shading caused by switching to a nontruthful mechanism such as a first price auction (FPA) rule is represented by the gap between the strategies for the first and second price bidding strategies in our dynamic setting. We also include the probability of winning in the bottom pane for reference.

We chose to plot bidding strategies under first- and second-price spot markets because they represent the polar extremes of static demand shading incentives among canonical pricing mechanisms. In that sense, the difference between the dash-dot line and the dashed line represents the maximal influence of static incentives for shaping behavior, and the difference between the solid line and the dash-dot line represents the influence of dynamic incentives. The conclusion we draw from the plot is that dynamic incentives tied to opportunity costs play a clearly dominant role in shaping behavior: for all bidder types with non-trivial win probabilities, the demand shading caused by intertemporal opportunity costs is an order of magnitude larger than the static demand shading.

Finally, our paper also makes a contribution to the theory underlying the large market models we use. The notion of a large market approximation, sometimes referred to as an Oblivious Equilibrium, is not novel to this paper. However, proving a formal relationship

between a model with a continuum of players and the finite markets that exist in reality is difficult when the market mechanism admits discontinuities, and an auction setting provides several points where such discontinuities can arise. We prove that despite these issues, one can view an equilibrium of the model with a continuum of buyers and sellers as an ε -equilibrium of the model with a finite number of buyers and sellers. We view this result as a justification for our use of the continuum model in our estimation and counterfactual exercises.

The remainder of this paper has the following structure. In Section 2 we develop a theory of bidding within a dynamic platform based on a model with a continuum of buyers and sellers. In Section 3 we use this model to specify a parsimonious structural model of eBay, which we show is identified from observables. We also propose a semi-nonparametric estimator based on B-splines. In Section 4 we present our model estimates. Section 5 presents our counterfactuals on welfare, the relative importance of platform composition and dynamic incentive effects, and optimal reserve prices. In Section 6, we prove that our model with a continuum of agents approximates an analogous model with a large, but finite, number of participants. Most of the proofs are relegated to the appendix.

1.1. Related Literature. The most closely related paper to ours is the contemporaneous Backus and Lewis [2012], which studies a model of eBay where bidders participate in a sequence of single-unit, second price auctions. Backus and Lewis [2012] focuses on identifying a model of buyer demand that includes product substitutability and the possibility that individual bidder demand evolves over time. In contrast, our focus is on using our model to measure the efficiency of the market platform, compute counterfactual studies of alternative market designs, and study the selection of buyers into the market. In addition, we provide a formal proof that the equilibria of our model with a continuum of the agents yields approximate equilibria of the analogous model with a finite set of buyers and sellers. Because of the very different focus of each work, we view our papers as complementary to one another.

Our methodology analyzes approximate equilibria played by a large number of agents, which has been a prominent theme in the microeconomics and industrial organization literatures. Due to the broad scope of this literature, we provide only a brief survey and a sample of the important papers related to the topic. Early papers focused on conditions under which underlying game-theoretic models could be used as strategic microfoundations for general equilibrium models (e.g., Hildenbrand [1974], Roberts and Postlewaite [1976], Otani and Sicilian [1990], Jackson and Manelli [1997]). Other early papers focused on conditions under which generic games played by a finite number of agent approach some limit game played by a continuum of agents (e.g., Green [1980],

Green [1984], and Sabourian [1990]). A more recent branch of this literature applies these ideas to simplify the analysis of large markets with an eye to real-world applications (e.g., Fudenberg, Levine, and Pesendorfer [1998]; MacLean and Postlewaite [2002]; Budish [2008]; Kojima and Pathak [2009]; Weintraub, Benkard, and Roy [2008]; Krishnamurthy, Johari, and Sundararajan [2014]; and Azevedo and Leshno [2016]).

Nekipelov [2007] and Hopenhayn and Saeedi [2016] develop models of intra-auction price dynamics with repeated bidding in a single auction. Their goal is to rationalize common empirical patterns concerning the timing of bids. In our model, we abstract away from intra-auction dynamics, and instead we concentrate on inter-auction dynamics and how future periods shape bidding incentives today. Peters and Severinov [2006] develop a model of a multi-unit auction environment similar to eBay with the goal of studying the sorting of buyers into sellers' auctions in a static setting.

Another related paper is Hickman [2010], which shows that the pricing rule on eBay is actually a hybrid of a first-price and a second-price mechanism due to the role of minimum bid increments. The sale price of an item on eBay is usually the second-highest bid plus a fixed increment, but when the top two bids are close enough (i.e., closer than the increment), then the sale price is set equal to the winner's bid. A rational eBay bidder accounts for the fact that her bid may affect the sale price, and the result is bids strictly between the first-price and second-price equilibria. Hickman, Hubbard, and Paarsch [2016] explore the empirical implications of the non-standard pricing rule on eBay within a static, one-shot auction model and show that estimates may become biased in an economically significant way if it is ignored. We build on these two papers in the following ways. First, our model incorporates both dynamic demand shading incentives (from future option value) and static demand shading incentives. Second, we extend the estimator of Hickman et al. [2016] to allow for binding reserve prices, which affects identification of the bidder arrival process and the private value distribution in complicated ways.

2. A MODEL OF PLATFORM MARKETS

Before describing our formal model of bidder behavior, we would like to informally describe the behavior our model is intended to capture. Our informal description is not meant to be universal, but we believe it to be fairly typical of behavior on eBay.

We imagine that before entering the eBay market, the buyer considers her own value for the good, makes an assessment of the opportunity cost, and formulates her bid. After entering the market, the bidder considers bidding in an auction that is closing in the near future. We assume that the time a bidder chooses to enter the market is driven by factors exogenous to eBay (e.g., the schedule of work breaks), which means means the

buyer only considers bidding on a small and randomly selected fraction of the auctions that close during that day. If a buyer wins the spot-market auction, then she does not participate in future days. All surviving buyers return to eBay the next day to place a bid. We summarize the timing in the following diagram.

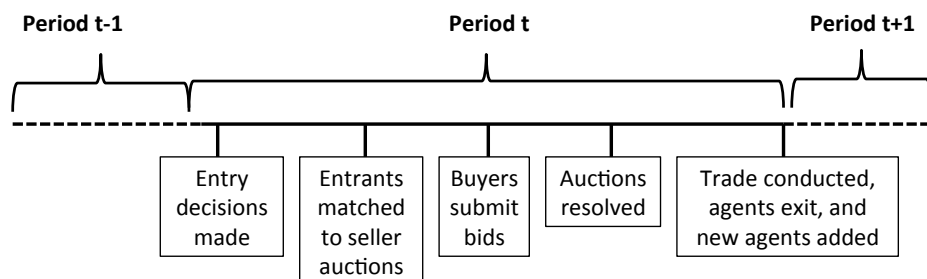


FIGURE 2. Timing within a period

Let us now address a few potential objections to our informal story. First, is it reasonable to assume that the buyer formulates her bid before entering the market? This behavior would be fully rational if eBay used a second price auction (SPA) format. However, since eBay auctions use a non-SPA pricing rule, there are static bid shading incentives that push the bidder to adjust her bid based on the competitive environment (e.g., the current highest bid) she observes upon entering an auction. However, as discussed in the introduction (see Figure 1) and in Section 4, intertemporal incentives that are independent of the number or valuations of the other bidders in the auction the bidder participates in today play a far more important role in determining the optimal bid than the competitive environment of a particular auction. Therefore the bidder has relatively weak incentives to update her bid according to the details or history of any given auction.

Second, by assuming a buyer participates in an auction closing near her time of entry, we rule out strategic entry into auctions. Strategic entry occurs to some extent - after all, there are likely to be at least a handful of auctions that close soon after the buyer enters. One might also worry that buyers select into auctions based on data unobservable to us that might be proxied for by, for example, the starting price of the auction (Roberts [2013]). If this were a significant issue, one would expect that there ought to be correlation between the starting price and the closing price of the auction. We find in our data that the correlation coefficient between the starting and the closing price is -0.015 and statistically indistinguishable from zero. We discuss this issue in more depth in Section 4.

Third, our informal story assumes the bidder does not repeatedly bid within a particular auction. As we discuss more thoroughly in Section 3, we only consider bids arriving

near the end of the auction. One reason we do this is to avoid the question of how to handle dynamic behavior within a given auction.⁵ As it stands, we see relatively few instances in the data of a bidder returning to increase her bid if she has already bid within the narrow window we consider at the close of each auction.

This section focuses on laying out the theoretical primitives of our continuum model, which is relatively simple to both estimate and solve for counterfactual equilibria. In Section 6 we prove that our model with a continuum of buyers and sellers is closely related to the analogous model with a finite number of buyers and sellers in the sense that a bidder’s dynamic value function under the continuum model approximates the value function arising from the finite model, and the approximation becomes arbitrarily accurate as the number of agents in the market increases. This implies that an exact equilibrium of the continuum model is an ε -equilibrium of the model with a finite number of agents.

Whenever possible, we develop theoretical results in terms that apply for arbitrary, well-behaved spot-market pricing rules so that our model may serve as a general framework for quantitative analysis of platform markets. Our definition of “well-behaved” is captured by Assumption 2.3 below.

2.1. Model Primitives. Now that we have described our model informally, we translate the story into a formal game-theoretic model. We treat sellers as a source of exogenous supply, and we take their decisions (e.g., entry/exit and starting prices) as exogenous and fixed within the model. This modeling choice is driven by the fact that sellers face very weak incentives to set the optimal starting price,⁶ which makes estimates based on a model of seller behavior less credible. Therefore, we present a theory of the buyer side of the market.

The market evolves in discrete time with periods indexed $t \in \{0, 1, 2, \dots\}$. In each period there is a measure 1 of sellers with reserve prices described by cumulative density function (CDF) G_R with support $[0, \bar{r}]$.^{7,8,9} G_R may have a mass point, but only at the lowest possible reserve price, $r = 0$, and has a probability density function (PDF)

⁵The question of intra-auction dynamics has been treated by Nekipelov [2007] and Hopenhayn and Saeedi [2016], and involves substantial complications that are beyond the scope of the current exercise. Other eBay models that view individual auction listings essentially as sealed-bid games include Bajari and Hortaçsu [2003], Hickman et al. [2016], Coey, Larsen, and Platt [2016], and Backus and Lewis [2016].

⁶As we discuss in greater length in section 5.3, the sellers earn less than \$1 in increased revenue by moving from a starting price of \$0 to the revenue maximizing starting price.

⁷Setting the measure of sellers to 1 is a normalization.

⁸The lowest starting price on eBay is \$0.99, but this does not affect our theoretical results.

⁹We use the letter G to refer to the CDFs of variables that are observable to the econometrician, and we reserve F to denote a CDF of an unobservable variable from the econometrician’s perspective.

$g_R(R|R > 0)$ that is strictly bounded away from zero over the rest of its support. We use the terms *starting price* and *reservation price* interchangeably.¹⁰

We refer to the set of buyers present at the start of period t as *potential entrants*; at the beginning of each period they make decisions based on the observed number and type distribution of the other potential entrants and their own types. Each period, the first choice each potential entrant must make is whether or not to enter the market and participate in the platform. We denote the choice to participate as *Enter* and refer to the agents that make this choice as *entrants*. The choice to not participate is denoted *Out*. If there is no history in which an agent chooses to *Enter*, then we assume that agent exits the game immediately and permanently. Otherwise any agent that chooses *Out* simply moves on to the next period.

Throughout we assume that the goods for sale are homogenous and that buyers have demand for a single unit.¹¹ Each buyer's value for the good is her private information, which we denote as v and assume is fixed over time. A buyer that wins a good on the eBay platform and pays a price of p receives a payoff in that period of

$$v_i - p - \kappa$$

where κ is a per-period bidding cost paid by entrants regardless of whether they win. We assume $\kappa > 0$; this may reflect the opportunity cost of time spent searching for a listing and participating in the market, or it may reflect an actual monetary participation fee that the platform designer imposes. Moving forward, we will use the terms "bidding cost" and "participation cost" interchangeably in reference to the parameter κ . If an entrant does not engage in trade, her payoff is simply $-\kappa$; a buyer that chooses not to enter the market earns a payoff of 0.

In period $t = 0$ there is a measure C_0 continuum of potential entrants with a type distribution equal to $F_{V,0}$.¹² The measure of potential entrants at the beginning of period t is denoted C_t . A measure of potential entrants equal to μ is added to the economy at the end of each period, and the distribution of the values of these new potential entrants has CDF $T_V(\cdot)$ with PDF $t_V(\cdot)$. We assume that t_V is strictly positive over the support $[0, 1]$. $F_{V,t}$ denotes the distribution the buyers' types in period t —including newly potential entrants and ones remaining from period $(t - 1)$ —and is an element of the space of probability measures over $[0, 1]$, denoted $\Delta([0, 1])$. Unless stated otherwise, $\Delta([0, 1])$ is

¹⁰eBay allows sellers to choose reservation prices that are hidden from buyers, but this is done so infrequently that we ignore it in our modeling.

¹¹We discuss the homogeneity of the goods in our data set in Section 4.

¹²Since the letter B is used later on to denote bids, we chose C , for "consumer," to represent the number of buyers in the market each period.

endowed with the weak-* topology. A generic, measure 0 buyer is denoted using the subscript i .

After choosing *Enter*, each entrant formulates a strategic bid without knowing either the number or identity of the other agents participating in the particular auction to which he or she is matched. The form of the random matching process, the distribution of entrant types, and the exogenous distribution of reserve prices is known to agents at the point when they choose their bids. We assume a simultaneous-move spot market; in other words, bidders maintain their ex-ante planned bid throughout the period and refrain from updating it during the life of their matched auction listing.

If a measure C_t of buyers choose to enter the auction market, the entrants are randomly assigned to auctions with each auction receiving a random number K of bidders where $Pr\{K = k\} = \pi(k, \lambda)$. For now, we impose no functional form on π , meaning the parameter vector $\lambda = \{\lambda_0, \lambda_1, \lambda_2, \dots\} \in \mathbb{R}^\infty$ is left unrestricted. We refer to λ as the *market tightness* parameter since it is determined by the buyer-seller ratio. The key properties we assume for the bidder arrival process are:

Assumption 2.1. *We require that π satisfy the following conditions:*

- (1) $E[K] = C_t$
- (2) π is continuous in C_t, F_V , and G_R
- (3) A local limit theorem applies, meaning that for the sequence (K_1, K_2, \dots) with $Z_N = \sum_{i=1}^N K_i$ and ψ denoting the density function of the normal distribution we have:

$$\sqrt{N\text{Var}[K]}Pr\{Z_N = k\} \rightarrow \psi\left(\frac{k - NE[K]}{\sqrt{N\text{Var}[K]}}\right) \text{ uniformly over } k \in \mathbb{Z} \quad (1)$$

The most novel assumption is (3), which requires that a local limit theorem apply. We use this assumption to approximate the probability mass function of sums of realizations of K using the probability density function of the normal distribution. Local limit theorems apply to most distributions of interest to economists including the generalized Poisson distribution used in our estimator.¹³ This level of generality will allow for a flexible empirical model specification later on.

Myerson [1998] showed that in games with stochastic participation, such as the spot market in our model, beliefs over the total number of competitors from the perspective of a participant in the auction (i.e., a bidder) are not the same as beliefs from the perspective of an outside observer (e.g., a seller or the platform designer). In particular, each bidder's beliefs about the number of other bidders in the auction to which she is matched are pinned down by the probabilities $\pi(\cdot; \lambda)$ but need not be the same. From the perspective

¹³See McDonald [2005] for more details and examples of local limit theorems.

of bidder 1, let M denote a random variable representing the number of competitors she faces, and let $\pi_M(M; \lambda)$ denote its probability mass function (PMF). As Myerson showed, π and π_M are the same distribution if and only if K is a Poisson random variable, a concept he referred to as *environmental equivalence*. Otherwise, her beliefs over M are given by

$$\pi_M(m; \lambda) = Pr [m \text{ opponents} | \lambda] = \pi(m+1; \lambda) \frac{(m+1)}{E[K]}. \quad (2)$$

Finally, we note at this point that our continuum model is “large” in the sense that the actions of individual bidders have no effect on the aggregate distribution of auction outcomes. However, the actions of individual bidders have a large effect on the auction to which they have been assigned. The tractability of the continuum model derives from the fact that, given knowledge of the equilibrium strategy, the distribution of types in the economy evolves deterministically.

2.2. Equilibrium. In this section we discuss the structure of the equilibrium under the assumption of a second-price auction rule since we can provide closed form solutions for some equilibrium quantities. As we show in Section 3.1.3 below, our general insights apply straightforwardly to other pricing mechanisms as well.

Bidding strategies can be written as functions $\mathcal{O} : [0, 1] \times \mathbb{R}_+ \times \Delta([0, 1]) \times \Delta([0, 1]) \rightarrow [0, 1]$ (\mathcal{O} for “offer”) with a typical bid denoted $\mathcal{O}(v, C_t, F_{V,t}, G_R)$. The entry decision for participating buyers is a function of the form $\theta : [0, 1] \times \mathbb{R}_+ \times \Delta([0, 1]) \times \Delta([0, 1]) \rightarrow \{Enter, Out\}$ with a typical realization $\theta(v, C_t, F_{V,t}, G_R)$. We let Σ denote the buyers’ strategy space.

We use the notation $x(b, C_t, F_{V,t}, G_R) = 1$ (0) to denote the random event that a buyer wins (loses) an auction with a bid of b , and $p(b, C, F_{V,t}, G_R)$ denotes the random transfer from the buyer to the seller/eBay conditional on a bid of b .¹⁴ To simplify notation we also define

$$\begin{aligned} \chi(b, C_t, F_{V,t}, G_R) &= E_t [x(b, C_t, F_{V,t}, G_R)] \\ \rho(b, C_t, F_{V,t}, G_R) &= E_t [p(b, C_t, F_{V,t}, G_R)] \end{aligned}$$

The expectation operator refers to the agent’s uncertainty regarding the other buyers that are participating in the auction to which he or she is matched. Note that ρ represents expected transfers that are not conditional on sale. That is, each entering bidder has an ex-ante expectation of paying $\rho(b)$ in the spot market, although under any winner-pay pricing rule only one bidder will pay a positive amount ex-post. For compactness we frequently suppress the notation for the aggregate state. We also often suppress the bid argument and assume the agent is following the equilibrium strategy.

¹⁴For example, $p(b, C, F_{V,t}, G_R) = 0$ if the buyer does not win the auction.

All agents discount future payoffs using a per-period discount factor $\delta \in (0, 1)$. The value function given a (symmetric) equilibrium strategy vector $\sigma = (\theta, \mathcal{O})$ for a bidder that chooses *Enter* is

$$\mathcal{V}(v, C_t, F_{V,t}, G_R | \sigma) = \chi v - \rho - \kappa + (1 - \chi) \delta \mathcal{V}(v, C_{t+1}, F_{V,t+1}, G_R | \sigma) \quad (3)$$

For a buyer that chooses *Out* we have

$$\mathcal{V}(v, C_t, F_{V,t}, G_R | \sigma) = \delta \mathcal{V}(v, C_{t+1}, F_{V,t+1}, G_R | \sigma) \quad (4)$$

We use the notation $\mathcal{V}(v, C_t, F_{V,t}, G_R | \sigma'_i, \sigma_{-i})$ when buyer i uses strategy σ'_i and all other agents follow σ .

Equilibrium requires that the actions taken are optimal with respect to the deterministic path of the state variables. We focus on stationary equilibria, which implies that the state variables are constant across time and the agent actions are optimal with respect to the state variables' realizations.

Definition 2.2. The strategy vector $\sigma = (\theta, \mathcal{O})$ and the states C and $F_V \in \Delta([0, 1])$ are a *Stationary Competitive Equilibrium* (SCE) if for all bidder values v we have

(1) For all $\sigma'_i \in \Sigma$,

$$\mathcal{V}(v, C, F_V, G_R | \sigma) \geq \mathcal{V}(v, C, F_V, G_R | \sigma'_i, \sigma_{-i})$$

(2) $\theta(v) = \text{Enter}$ if and only if

$$\chi v - \rho - \kappa + (1 - \chi) \delta \mathcal{V}(v, C, F_V, G_R | \sigma) \geq 0$$

(3) $C = C_t = C_{t+1}$ and $F_V = F_{V,t} = F_{V,t+1}$ are consistent with the laws of motion of the game.

In an SCE, the agents bid the same amount in each period, meaning that the bidding function can be written $\beta : [0, 1] \rightarrow [0, 1]$. The entry decision must take the form

$$\theta(v, C, F_V, G_R) = \text{Enter} \text{ if and only if } \chi [v - \delta \mathcal{V}(v, C, F_V, G_R | \sigma)] - \rho \geq \kappa$$

Any buyer that is indifferent between entering and staying out must have a continuation value of 0 since, due to the stationarity, if she is indifferent today she will be indifferent in every future period. This implies that new entrants will either exit immediately or enter the market in every period. Because of this structure, we can describe the entry strategies through a cutoff function $e(C, F_V, G_R) = \inf_v \{v : \chi v - \rho \geq \kappa\}$ and

$$\theta(v, C, F_V, G_R) = \text{Enter} \text{ if and only if } v \geq e(C, F_V, G_R) = \underline{v} \quad (5)$$

where \underline{v} is the lowest value buyer willing to enter the market.

In an SCE under a SPA spot price mechanism, it is an equilibrium in weakly undominated strategies for a bidder to bid his value for the good minus the opportunity cost of

winning. In the static, one-shot setting, this opportunity cost is 0 since outside options are assumed not to exist. In our dynamic model, the opportunity cost of winning today is the continuation value the bidder receives if she instead returns to the market to bid again in a future period. Therefore we can write

$$\beta(v) = v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma) \quad (6)$$

In a static auction, one's equilibrium bid is chosen to balance out opposing forces: a higher bid will increase the chance of winning, but it may also raise the price one pays as well. Whenever the second force is present, bidders shade their demand. Henceforth, we refer to these two forces as *static incentives*. Intertemporal dynamics introduce an additional incentive for bidders to shade their bids: if the spot-market price is sufficiently high today, then a bidder would prefer to wait in expectation of lower prices tomorrow. Therefore, even when the spot-market game follows a second-price rule, rational bidders in equilibrium engage in demand shading. In what follows, we refer to this source of demand shading as *dynamic incentives*. Since a buyer with value \underline{v} is indifferent between entering and staying out of the market, we must have $\mathcal{V}(\underline{v}, C, F_V, G_R|\sigma) = 0$, which in turn means the lowest value type that enters does not shade her bid (i.e., $\beta(\underline{v}) = \underline{v}$).

At this point, let us take a moment to define what it means for the state variables C and F_V to be consistent with the laws of motion of the game. In other words, what are the conditions an SCE must satisfy for stationarity to hold? These conditions are used implicitly in both our estimator and our counterfactual calculations. Before we begin, note that the structural primitives of our model are μ , κ , T_V , and G_R , so our stationarity conditions will be written in terms of these quantities as well as the equilibrium strategy.

For an economy to be stationary, the distribution and measure of buyers that win auctions and exit the game must be replaced by an identical distribution of new entrants. For this to be true, we must have:

$$\text{For all } v, \mu t_V(v) = \chi(\beta(v)) f_V(v) \mathcal{C} \quad (7)$$

As before, $\beta(v)$ is the symmetric equilibrium bidding strategy, $\chi(b)$ is the probability of winning a spot-market auction with a bid of $\beta(v)$, and \mathcal{C} is the market-wide buyer-seller ratio. The left-hand side of Equation 7 is the measure of buyers of type v entering the market, and the right-hand side is the density of buyers of type v who win an auction and exit the market. Equation 7 and the fact that $F_V(1) = 1$ pins down f_V and \mathcal{C} .

This raises the question of how one can compute β and χ . For general pricing rules, β can be formulated as the solution to an ordinary differential equation determined by the first order conditions of the problem facing a bidder. In an SPA setting, we have $\beta(v) = v - \delta\mathcal{V}(v)$. To compute \mathcal{V} , we use Theorem 1 of Pavan, Segal, and Toikka [2014],

which implies that \mathcal{V} can be described as:

$$\mathcal{V}(v) = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \int_{\underline{v}}^v (1 - \chi(\beta(s)))^{\tau-t} \chi(\beta(s)) ds \quad (8)$$

and we can write χ as

$$\chi(b) = G_R(b) \sum_{m=0}^{\infty} \pi_M(m) F_B(\beta^{-1}(b))^m \quad (9)$$

Given Equations 7 - 9 we can pin down all of the endogenous variables from the structural primitives.

In the interest of generality, we extend our framework to account for a variety of payment rules. For our large-market approximation results (Section 6) to hold for platforms with non-second-price spot market mechanisms, we require the following assumption on the SCE strategy:

Assumption 2.3. *The best response bids of the buyers are continuous in the sup-norm with respect to the parameter C and the distribution of bids of the entrants as long as the distribution of bids admits a PDF that is bounded from above.*

Fix a distribution of bids that admit a PDF g_B bounded from above. We assume that we can choose some $\varphi \in (0, 1)$ such that for any best response by the buyers, denoted b , to (C, g_B, G_R) and any $v > v'$ we have

$$b(v) - b(v') \in \left[\varphi(v - v'), \frac{v - v'}{\varphi} \right] \quad (10)$$

For the duration of this paper, we will take Assumption 2.3 as given.^{15,16}

In summary we can describe any SCE as a vector of strategies, $\sigma = (e, \beta)$ and aggregate states (C, F_V, G_R) . Our next result (see online technical appendix for proof) shows that there exists an SCE of our model.

Proposition 2.4. *A stationary competitive equilibrium exists, and a positive mass of buyers choose to enter the market if κ is not too large.*

The main difficulties in the proof are (1) arguing we can limit consideration to a compact strategy and state space and (2) ruling out a number of utility discontinuities that naturally arise in auction markets. Once we handle these issues, our proof uses a traditional fixed point argument.

¹⁵Given Theorem 6.2, we could have equivalently required continuity of the bidding equilibrium of the static game ($\delta = 0$).

¹⁶In the SPA we have $\beta(v) = v - \delta\mathcal{V}$. Since the value function must be strictly increasing with a slope less than 1, we know that $\beta(v)$ satisfies equation 10. Therefore the SPA satisfies assumption 2.3.

3. AN EMPIRICAL MODEL OF DYNAMIC PLATFORM MARKET BIDDING

We now shift focus to developing a structural model based on the SCE described above. Letting L denote sample size (where an auction is the unit of observation), the observables, $\{\tilde{k}_l, r_l, y_l\}_{l=1}^L$, include \tilde{k}_l , the observed number of bidders within the l^{th} auction; r_l , the reserve price; and y_l , the highest losing bid. For the purpose of our discussion on identification, we leave the bidder arrival process $\pi(\cdot; \lambda)$ nonparametric, so that the market tightness parameter vector λ is allowed to be infinite-dimensional with $\lambda_k = \Pr[K = k]$, $k = 0, 1, 2, \dots$. For convenience we drop the parameter argument in the PMF $\pi(\cdot)$ unless context requires specificity.

For simplicity of discussion, consider the decision problem of a bidder who has decided to enter and finds herself competing within a spot-market auction; we will refer to her as bidder 1. As before, denote the total number of opponents she faces by $M \equiv K - 1 \geq 0$ and recall that from 1's perspective $\pi_M(\cdot)$ may not be the same distribution as $\pi(\cdot)$, though λ determines both. Prior to bidding, bidder 1 observes her own private valuation v and she views her opponents' private values as independent realizations of a random variable $V \sim F_V$ having strictly positive density f_V on support $[\underline{v}, \bar{v}]$ with $\underline{v} > 0$.¹⁷ The theory from the previous section depicted a set of potential buyers, some of whom choose to enter the bidding market and some of whom don't, with Equation 5 determining the relevant entry cutoff. Since we are unable to collect real-world observations on non-entrants, we shift notation slightly from the previous section and adopt the convention that F_V is the steady-state distribution of buyer types who choose *Enter*. This is possible because, in an SCE, whenever it is optimal for a buyer to enter (stay out) in a given period, it will always be optimal for her to enter (stay out) in every future period until she wins an auction and exits. Recall that \underline{v} is the type that is just indifferent to entering, leading to the following formula that we refer to as the "zero surplus condition":

Assumption 3.1. $\mathcal{V}(\underline{v}) = 0$

Bidder 1 wishes to formulate an optimal bid that reflects her static incentives (from competition in the spot market) and her dynamic incentives (from the option value of re-entering the market in the future if she loses today). She views the bids of her opponents as a random variable $B = \beta(V) \sim G_B(B) = F_V[\beta^{-1}(B)]$ with support $[\underline{b}, \bar{b}]$. Let B_M denote the maximal bid among all of bidder 1's opponents with a distribution defined by $G_{B_M}(B_M) = \sum_{m=1}^{\infty} \frac{\pi_M(m)}{1 - \pi_M(0)} G_B(B_M)^m$. R is the starting price of the auction, which is

¹⁷Nothing in our theory relied on values being drawn from the specific $[0, 1]$ interval, so it is innocuous to have values drawn from some other compact interval of real numbers.

randomly drawn from CDF G_R . In order to win, player 1's bid must exceed the realized value of the random variable:

$$Z \equiv \begin{cases} R & \text{if } M = 0 \\ \max\{R, B_M\} & \text{otherwise} \end{cases}$$

Note that the distribution of Z is the same as the win probability function:

$$\chi(b) = G_R(b) \sum_{m=0}^{\infty} \pi_M(m) G_B(b)^m \quad (11)$$

3.1. Model Identification. We first establish nonparametric identification results in Sections 3.1.1–3.1.2 for a baseline case with a second-price spot-market auction. In Section 3.1.3 we extend our identification result to the case where the spot-market game is non-second price, which creates additional static demand shading incentives. This extension will be useful in dealing with data from eBay, which employs a hybrid pricing mechanism that exhibits elements of both first-price and second-price rules.

3.1.1. Baseline Model: Second-Price, Sealed-Bid Spot-Market Auctions. A second-price spot-market mechanism implies a specific form for the expected payment function,

$$\begin{aligned} \rho(b) \equiv \mathbb{E}[p_B(b)] &= \pi_M(0) G_R(b) \mathbb{E}[R | R \leq b] \\ &+ [1 - \pi_M(0)] \int_b^b t [g_R(t) G_{B_M}(t) + G_R(t) g_{B_M}(t)] dt \end{aligned} \quad (12)$$

Since the market is in steady-state, we can express the Bellman equation and bidding strategy as

$$\mathcal{V}(v) = \max_{b \in \mathbb{R}_+} \left\{ \chi(b)v - \rho(b) - \kappa + [1 - \chi(b)] \delta \mathcal{V}(v) \right\} \quad (13)$$

$$\beta(v) = v - \delta \mathcal{V}(v) \quad (14)$$

The demand shading factor given by bidder 1's continuation value, $\delta \mathcal{V}(v)$, is uniquely characterized by four things: the per-period entry cost, κ ; the distribution of bids, $G_B(b)$; the distribution of starting prices, G_R ; and the market tightness parameters, λ , that determine the overall ratio of buyers and sellers. Thus, mapping bids into private values requires first identifying these four objects. The model is said to be identified if there exists a unique set of structural primitives that could rationalize a given realization of the joint distribution of observables, $\{\tilde{k}_l, r_l, y_l\}_{l=1}^L$. The structural primitives to identify are μ, κ, T_V , and G_R .¹⁸

¹⁸Our exposition might occasionally appear to treat F_V as a primitive, but in reality it is pinned down by Equation 7. As the reader will see, we use a version of Equation 7 to estimate T_V from F_V (and other variables).

Note that (14) implies $\mathcal{V}(v) = \frac{v - \beta(v)}{\delta}$. By substituting this expression into equation (13) and using the shorthand notation $b^* = \beta(v)$, we can rearrange terms to get

$$v = b^* \frac{1 - \delta(1 - \chi(b^*))}{1 - \delta} - \frac{\delta}{1 - \delta} (\rho(b^*) + \kappa) \equiv \beta^{-1}(b^*) \quad (15)$$

The expression for the inverse bidding function above is crucial to demonstrating identification of the model, as it shows that the relationship between bids and latent private values is fully characterized by a small set of parameters and functionals. Specifically, if the econometrician can identify λ , κ , G_R , and G_B from the observables, then for a given discount factor δ we can reverse engineer the value v that rationalizes bid b as a best response to prevailing market conditions.

3.1.2. Identification of the Bid Distribution and Bidder Arrival Process. One challenge to empirical work is that the observed number of bidders in each auction, \tilde{K} , is only a lower bound on the actual number of bidders matched to the auction, K . Due to random ordering of bid submission times across all bidders who watch an item with intent to compete, some may find that their planned bid was surpassed before they had a chance to submit it to the online server. These bidders will never be visible to the econometrician, even though they were matched to and competing to win the auction.

To solve this problem we incorporate an explicit model of the sample selection process into our identification strategy. In doing so we adopt an approach similar to that of Hickman et al. [2016] who proposed a model of a *filter process* executed by Nature that randomly withholds some bidders from the econometrician's view.¹⁹ For a given auction with k total matched bidders, this filter process first randomly assigns each bidder an index $\{1, 2, \dots, k\}$, where one's position in the list determines the ordering of bid submission times. Nature then visits each bidder in the order of her index within the list, keeping a running record of the current lead bidder and current price as she goes. As Nature visits each bidder in the list, she only records bid tenders that cause her running record of the price or lead bidder to update (i.e., those that exceed the second highest from among previous bid tenders). Otherwise, Nature skips bidder i 's submission as if it never happened and reports to the econometrician only the record of price path updates, which reveals $\tilde{k} \leq k$ observed bidder identities. This filter process is meant to depict the way information is recorded on real-world platform markets like eBay, and it opens up the possibility that some bidders will not appear to have participated even though they had an intent to bid. This view of intra-auction dynamics assumes that the ordering of bidders' submission times is random. Note that we remain agnostic on how

¹⁹In a similar setting, Platt [2015] explored parametric inference assuming that K is Poisson distributed.

agents decide to bid early or late; we merely rule out the possibility of bidder collusion on the ordering of their submissions.

Since the filter process does not depend on the particular distribution of K , the distribution of \tilde{K} conditional on a given k can be characterized without knowing λ . Moreover, since a bidder's visibility to the econometrician only depends on whether her bid exceeds the second-highest preceding bid, the researcher can easily simulate the filter process based on quantile ranks—without knowing G_B or F_V a priori—to compute conditional probabilities $\Pr[\tilde{k}|k]$ for various (\tilde{k}, k) pairs.²⁰ We adopt a special notation for this object, $P_0(\tilde{k}, k) \equiv \Pr[\tilde{k}|k]$ and treat it as an observable. Since \tilde{k} is observable, we can use this information to express its PMF, denoted $\tilde{\pi}(\tilde{k})$, as a function of the market tightness parameters λ : $\tilde{\pi}(\tilde{k}) = \sum_{k=\tilde{k}}^{\infty} P_0(\tilde{k}, k)\pi(k; \lambda)$.

However, this equation will not suffice as a basis for identification and estimation in our case. Unlike Hickman et al., our empirical application requires us to allow for the presence of binding reserve prices. These introduce a second layer of selection, driving a further wedge between actual participation k and observed participation, \tilde{k} . Not only do some bidders go unobserved because the filter process withholds them from view, but an additional fraction of bidders, who would have otherwise been reported by Nature, go unobserved because their bids fall below the reserve price. This second layer of selection produces substantial complications since G_B now determines how the second source of selection influences the relation between the distribution of observed \tilde{K} and the underlying distribution of actual K .

In order to solve this problem we propose an adjusted filter process wherein, for each auction, Nature randomly draws k from $\pi(k)$, r from G_R , and an ordered list of bidders with timing index, $i \in \{1, 2, \dots, k\}$. Each bidder is endowed with an iid private value v_i drawn from F_V . The bidders formulate their strategic bids without knowing the realization of k or r , and Nature then compiles a reported list of bidders for the econometrician in two steps. First, she visits each bidder in the list and dismisses anyone whose strategic bid does not meet the reserve price r . Second, Nature assigns the remaining set of $k' \leq k$ bidders new indices $i' \in \{1, 2, \dots, k'\}$ in increasing order of their raw indices i , and then executes the standard filter process algorithm for computing and reporting \tilde{k} conditional on r . Finally, Nature reports \tilde{k} and r to the econometrician.

In order to characterize the conditional distribution of \tilde{K} given r , first note that if there are $K = k$ total bidders, the probability that exactly j of them are screened out by r is

²⁰Hickman et al. [2016] simulated 10^{12} auction filter processes to obtain a lower-diagonal matrix of conditional probabilities $\Pr[\tilde{k}|k]$, for each $\tilde{k} \leq k$ and $k \leq 100$. With that many simulations, the element-wise approximation error is on the order of $\sqrt{10^{-12}} = 10^{-6}$, and their simulated matrix can be re-used for any setting in which $E[K] \leq 40$.

$\binom{k}{j} G_B(r)^j [1 - G_B(r)]^{k-j}$. Now suppose there are \tilde{K} observed bidders in an auction with N total bidders. We can combine the two levels of selection in the adjusted filter process with the following equation:

$$\Pr[\tilde{K} = \tilde{k} | K = k, r] = \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} G_B(r)^j [1 - G_B(r)]^{k-j} P_0(\tilde{k}, k - j) \quad (16)$$

The sum is to account for the fact that any number of bidders between 0 and $k - \tilde{k}$ could be screened out by selection on reserve prices. The trailing term on the end is to account for the standard filter process running its course with the surviving set of bidders. Equation (16) now allows us to characterize the distribution of observed \tilde{K} conditional on the observable reserve price r , as

$$\tilde{\pi}(\tilde{k} | r) = \sum_{k=\tilde{k}}^{\infty} \Pr[\tilde{k} | k, r] \pi(k; \lambda). \quad (17)$$

Estimation of λ can no longer be separated from G_B because Equation 17 involves both of these objects. Fortunately though, this is merely a matter of implementation, as the following demonstrates that the model is nonparametrically identified from the available observables.

Proposition 3.2. *For a given discount factor δ , the market tightness parameters λ , bidding cost κ , and steady-state measures F_V , μ , and T_V are nonparametrically identified from the joint distribution of the observables $\{\tilde{k}_l, r_l, y_l\}_{l=1}^L$ when the spot market mechanism is a sealed-bid, second-price auction.*

Proof. $H(\cdot)$ denotes the distribution of the highest losing bid from the econometrician's perspective, and $H(\cdot)$ takes the form

$$H(b) = \sum_{k=2}^{\infty} \frac{\pi(k; \lambda)}{1 - \pi(0; \lambda) - \pi(1; \lambda)} \left(G_B(b)^k + k G_B(b)^{k-1} [1 - G_B(b)] \right). \quad (18)$$

$H(b)$ is a weighted average of the distributions of second order statistics from samples of varying k , where the weights are the probability that a given k will occur as the number of bidders matched to a particular listing. If we let $\varphi(H(b); \lambda) = G_B(b)$ denote the inverse of (18), then it follows that, holding λ fixed, φ is monotone in $H(b)$ for each b . Combined with the fact that $H(b)$ and $\tilde{\pi}(k | r)$ are known, this implies that λ and G_B are identified from the observables and from equations (17) and (18). Moreover, G_R is directly observable from data.

Given these three pieces, it also follows that the win probability $\chi(b)$ and the expected winner payment $\rho(b)$ are identified through equations (11) and (12) above. To identify

the participation cost, combine equation (13) with the zero surplus condition (3.1) to find the following relation:

$$\chi(\underline{v})\underline{v} - \rho(\underline{v}) = \kappa \quad (19)$$

In other words, the marginal market participant reaps just enough benefit in expectation to offset the cost of participation.

With $\chi(b)$, $\rho(b)$, and κ known, equation (15) shows that β^{-1} is also identified if the discount factor δ is known, and in turn, the private value distribution is identified through the relationship $F_V(v) = G_B[\beta(v)]$. With F_V known, μ is identified through either of the following two equivalent expressions which determine the mass of transactions each period, and therefore the total mass of buyers exiting the market:

$$\begin{aligned} \mu &= \int_{\underline{v}}^{\bar{v}} \chi[\beta(v)] f_V(v) dv \\ &= [1 - \pi(0)] G_R(\underline{b}) + \int_{\underline{b}}^{\bar{b}} g_R(r) \left(\sum_{k=1}^{\infty} \pi(k) [1 - G_B(r)^k] \right) dr \end{aligned} \quad (20)$$

Finally, once μ is known T_V is identified through equation (7). \square

3.1.3. Model Identification Under Alternative Spot Market Mechanisms. We now extend our identification result to cover platform markets that use alternative spot-market pricing mechanisms. Under the second-price platform model above, we saw that market dynamics produce an incentive to engage in demand shading due to the option value of future market participation in the event of a loss. Alternative spot-market mechanisms in which the winner's bid directly influences the current-period sale price will produce further incentives for demand shading. As we show below, this static demand shading margin is layered on top of the dynamic demand shading from the baseline second-price model in an intuitive way.

To make these ideas more concrete, we continue to use $\rho(b)$ to denote the expected payment under the prevailing spot market mechanism, whatever it may be. Conditional on choosing to enter the platform market, once a bidder is matched to a seller her decision problem is:

$$\beta(v) = \arg \max_b \left\{ \chi(b)v - \rho(b) - \kappa + [1 - \chi(b)]\delta\mathcal{V}(v) \right\}$$

We find it useful to refer to a bidder's private value minus her opportunity cost as her *dynamic value*, denoted $\tilde{v}_v \equiv v - \delta\mathcal{V}(v)$. By rearranging terms we can re-cast the bidder's decision problem as choosing a functional $\tilde{\beta} : \tilde{V}_V \rightarrow \mathbb{R}_+$ to optimize:

$$\tilde{\beta}(\tilde{v}_v) = \arg \max_b \left\{ \chi(b)\tilde{v}_v - \rho(b) + c \right\}, \quad (21)$$

where $c = -\kappa + \delta\mathcal{V}(v)$ is a constant and $\beta(v) = \tilde{\beta}(\tilde{v}_v)$. In other words, under alternative spot-market pricing rules agents shade demand as if they were in a static one-shot auction, but where shading is relative to their dynamic value \tilde{v}_v (see Proposition 6.2). This parsimonious layering of demand shading incentives is useful because it allows us to show that if the right set of observables are available to identify the mapping $\tilde{\beta}$ that would arise in a static, one-shot auction with allocation rule χ and pricing rule ρ , then the value function \mathcal{V} and the private value v from the dynamic auction market is also identified. To see why, note that by plugging the optimizer $\tilde{\beta}$ into equation (13) and rearranging we get:

$$\mathcal{V}(v) = \frac{\chi[\tilde{\beta}(\tilde{v}_v)]v - \rho[\tilde{\beta}(\tilde{v}_v)] - \kappa}{1 - \delta(1 - \chi[\tilde{\beta}(\tilde{v}_v)])}$$

Using the shorthand $b^* = \tilde{\beta}(\tilde{v}_v) = \beta(v)$ and substituting in the definition of \tilde{v}_v , we can rearrange terms further to get:

$$v = \tilde{v}_v \left(\frac{1 - \delta[1 - \chi(b^*)]}{1 - \delta} \right) - \frac{\delta}{1 - \delta} (\rho(b^*) + \kappa) = \beta^{-1}(b^*). \quad (22)$$

In the case of a SPA spot market, where $b^* = \tilde{\beta}(\tilde{v}_v) = \tilde{v}$, equation (22) reduces to equation (15) above.

Proposition 3.3. *For a given discount factor δ , the market tightness parameters λ , bidding cost κ , and steady-state measures F_V , μ , and T_V are nonparametrically identified under any spot market mechanism for which either*

- (1) *the optimizer of (21) is scalar-valued and the allocation rule $\chi(b)$ and pricing rule $\rho(b)$ can be identified from the available observables $\{\tilde{k}_l, r_l, y_l\}_{l=1}^L$; OR*
- (2) *the optimizer of (21) could be identified from the available observables $\{\tilde{k}_l, r_l, y_l\}_{l=1}^L$ if they were generated from a sample of static, one-shot auction games.*

Proof. The argument for identification of λ , G_B , and G_R is the same as in Proposition 3.2. For case (1), assuming that $\chi(b)$ and $\rho(b)$ can be expressed as a function of observable objects (including λ , G_B , and G_R), equations (19) and (21) identify κ , $\tilde{\beta}(\cdot)$, and \tilde{v}_v . For case (2), consider a hypothetical alternative world where the same set of observables were actually generated from a sample of static, one-shot auctions, based on underlying private valuations \tilde{v}_v . If the observables (including λ , G_B , and G_R) are known to identify the inverse bid mapping in that static world, then once again we can treat κ , $\tilde{\beta}(\cdot)$, and \tilde{v}_v as known.

Finally, equation (22) maps each observed bid b into a private value v that rationalizes b as a best response to market conditions both within-period and future. This implies

that F_V is identified, after which equations (7) and (20) identify μ and T_V similarly as before. \square

Proposition 3.3 is useful because it broadens the applicability of our model and methodology to allow for empirical work for any spot-market mechanism that admits a monotone equilibrium in the static setting and for which the pricing and allocation rules can be expressed in terms of λ , G_R , and G_B . For example, any platform model where the spot market uses a first-price rule will still be nonparametrically identified given commonly available observables. The structural auctions literature has established a broad array of nonparametric identification results for settings of static, one-shot auctions, beginning with the work of Guerre, Perrigne, and Vuong [2000] and Athey and Haile [2002]. The result above allows for the researcher in a dynamic marketplace to use established, static-market identification strategies in a variety of settings, provided they can be adapted to handle stochastic participation with a known matching process $\pi(\cdot; \lambda)$. The ability to incorporate established identification strategies for static auctions will be useful as we develop an estimator for eBay data. There, the pricing rule is known to be a non-standard combination of both first-price and second-price rules, which causes bidders to engage in additional demand shading from their static, strategic incentives.

3.2. A Two-Stage, Semi-Parametric Estimator. Thus far in our discussion we have left the bidder arrival process $\pi(k; \lambda)$ unrestricted in order to demonstrate that the theoretical model is sufficient on its own (given our observables) to identify the structural primitives without resorting to parametric assumptions. In this section we develop an estimator to implement our identification strategy, but for the sake of tractability we now assume K follows a generalized Poisson distribution (Consul and Jain [1973]) with PMF:

$$\pi(K = k; \lambda) = Pr [K = k | \lambda] = \lambda_1 (\lambda_1 + k\lambda_2)^{k-1} \frac{e^{-(\lambda_1 + k\lambda_2)}}{k!}, \quad \lambda_1 > 0, \quad |\lambda_2| < 1, \quad (23)$$

The first two moments of the generalized Poisson distribution are $E[K] = \lambda_1 / (1 - \lambda_2)$ and $\text{Var}[K] = E[K] / (1 - \lambda_2)^2$. While the generalized Poisson reduces to a regular Poisson distribution when $\lambda_2 = 0$, it exhibits fatter tails when $\lambda_2 > 0$ and thinner tails when $\lambda_2 < 0$. Given the linkage between the traditional and generalized Poisson distributions, we refer to λ_1 as the *size parameter* and λ_2 as the *dispersion parameter*. Developing an estimator based on finite-dimensional λ avoids significant complications that we discuss briefly below, but which are beyond the scope of this work.

Recall from the classic Myerson [1998] result that we have environmental equivalence—that is, a participant’s beliefs over the total number of competitors in an auction correspond to those of an outside observer—only in the special case of Poisson-distributed K . In general, bidder 1’s beliefs about the number of her opponents, M , follows $\pi_M(m, \lambda) =$

$\pi(m+1; \lambda)(m+1) \frac{(1-\lambda_2)}{\lambda_1}$. Since the generalized Poisson with $\lambda_2 > 0$ (< 0) admits an unusually high (low) number of large auctions relative to the standard Poisson distribution, each bidder believes that, conditional on herself having been matched into an auction, it is likely that it will be one with many (few) other bidders. It is easy to confirm that by plugging in $\lambda_2 = 0$ participant beliefs π_M become Poisson like outsider beliefs π .

Following our identification argument, G_B and λ must be jointly estimated, which rules out many common methods such as kernel smoothing. For our purpose, we opt for the method of sieves approach (see Chen [2007]) where a finite-dimensional, parametric form is imposed on G_B in finite samples and made to be ever more flexible as the sample size increases. We choose to specify G_B as a B-spline, which is a linear combination of globally defined basis functions that mimic the behavior of piecewise, local splines (the name ‘‘B-splines’’ is short for *basis splines*). By the Stone–Weierstrass theorem, B-splines can be used to approximate any continuous function to arbitrary precision given sufficiently many basis functions.²¹ B-splines provide a remarkable combination of flexibility and numerical convenience that is ideally suited to our application.

Let $\mathbf{n}_b = \{n_{b1} < n_{b2} < \dots < n_{b, I_b+1}\}$ be a set of knots on bid domain $[\hat{b}, \hat{b}] = [\min_l \{y_l\}, \max_l \{y_l\}]$ that create a partition of I_b subintervals. This need not be a uniform partition, but we do require that $n_{b1} = \hat{b}$ and $n_{b, I_b} = \hat{b}$ so that the partition spans the entire domain space. The knot vector, in combination with the Cox-de Boor recursion formula, uniquely defines a set of $I_b + 3$ cubic B-spline basis functions $\mathcal{F}_{b,i} : [\hat{b}, \hat{b}] \rightarrow \mathbb{R}$, $i = 1, \dots, I_b + 3$ that give us our parameterization of the bid distribution:²²

$$\hat{G}_B(b; \alpha_b) = \sum_{i=1}^{I_b+3} \alpha_{b,i} \mathcal{F}_{b,i}(b)$$

We also follow this approach for estimating G_R and F_V . Let $\mathbf{n}_r = \{n_{r1} < n_{r2} < \dots < n_{r, I_r+1}\}$ and $\mathbf{n}_v = \{n_{v1} < n_{v2} < \dots < n_{v, I_v+1}\}$ denote knot vectors for the reserve price distribution and private value distribution, defining I_r and I_v subintervals, respectively. The former is chosen to span $[\underline{r}, \hat{r}] = [0.99, \max_l \{r_l\}]$ and the latter spans $[\hat{v}, \hat{v}]$, with the bounds to be estimated. These knot vectors determine our other basis functions $\mathcal{F}_{r,i} : [\underline{r}, \hat{r}] \rightarrow \mathbb{R}$, $i = 1, \dots, I_r + 3$ and $\mathcal{F}_{v,i} : [\hat{v}, \hat{v}] \rightarrow \mathbb{R}$, $i = 1, \dots, I_v + 3$ which in turn render our parameterizations $\hat{G}_R(r; \alpha_r) = \sum_{i=1}^{I_r+3} \alpha_{r,i} \mathcal{F}_{r,i}(r)$ and $\hat{F}_V(v; \alpha_v) = \sum_{i=1}^{I_v+3} \alpha_{v,i} \mathcal{F}_{v,i}(v)$.

²¹Unlike global polynomials (e.g., Chebyshev), B-splines are capable of accommodating an unbounded degree of curvature at a point with finitely many terms if the researcher has a priori information on regions of the functional domain where such flexibility is needed.

²²A standard text on B-splines is de Boor [2001]. See also [Hickman et al., 2016, Online Appendix] for a brief but detailed primer on construction of B-spline basis functions, their derivatives, and their advantages for empirical work in economics.

Following our identification argument, we separate estimation into two stages. In the first stage we flexibly estimate λ , G_B , and G_R , and in the second stage we construct the remaining objects $\chi(\cdot)$, $\rho(\cdot)$, κ , $\tilde{\beta}^{-1}(\cdot)$, $\beta^{-1}(\cdot)$, $\mathcal{V}(\cdot)$, $F_V(\cdot)$, μ , and T_V as functions of first-stage parameter estimates. Note that Stages I and II differ in that Stage I is an estimation step, but Stage II is a purely computational step based on the outputs from Stage I.

3.2.1. *Stage I: λ , G_B , and G_R .* Recalling that the matrix of conditional probabilities $P_0(\tilde{k}, k)$ is known beforehand, in a slight adjustment of notation we now define the model-generated conditional PMF of \tilde{K} given r as

$$\tilde{\pi}(\tilde{k}|r; \lambda, \alpha_b) = \sum_{k=\tilde{k}}^{\bar{K}} \left\{ \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} \hat{G}_B(r; \alpha_b)^j [1 - \hat{G}_B(r; \alpha_b)]^{k-j} P_0(\tilde{k}, k-j) \right\} \pi(k; \lambda) \quad (24)$$

where \bar{K} is an upper bound on the auction sizes we consider. We also adopt the following as the empirical analog of the conditional PMF:

$$\hat{\pi}(\tilde{k}|r) = \sum_{l=1}^L \mathbb{1}(\tilde{k}_l = \tilde{k}) \frac{\mathcal{K}\left(\frac{r-r_l}{h_R}\right)}{\sum_{t=1}^L \mathcal{K}\left(\frac{r-r_t}{h_R}\right)} \quad (25)$$

where $\mathbb{1}(\cdot)$ is an indicator function, \mathcal{K} is a boundary-corrected kernel function, and h_R is an appropriately chosen bandwidth.²³ Finally, we define the model-generated highest loser bid distribution as

$$H(b; \lambda, \alpha_b) = \sum_{k=2}^{\infty} \frac{\pi(k; \lambda) (G_B(b; \alpha_b)^k + kG_B(b; \alpha_b)^{k-1} [1 - G_B(b; \alpha_b)])}{1 - \pi(0; \lambda) - \pi(1; \lambda)}, \quad (26)$$

and its empirical analog as $\hat{H}(b) = \sum_{l=1}^L \mathbb{1}(y_l \leq b) / L$. Using these separate pieces we can define a method of moments estimator as

$$\begin{aligned} (\hat{\lambda}, \hat{\alpha}_b) &= \arg \min_{(\lambda, \alpha_b) \in \mathbb{R}^{I_b+5}} \sum_{l=1}^L \left\{ [\tilde{\pi}(\tilde{k}_l|r_l; \lambda, \alpha_b) - \hat{\pi}(\tilde{k}_l|r_l)]^2 + [H(y_l; \lambda, \alpha_b) - \hat{H}(y_l)]^2 \right\} \\ &\text{subject to} \\ &\alpha_{b,1} = 0, \quad \alpha_{b, I_b+3} = 1, \\ &\alpha_{b,i} \leq \alpha_{b,i+1}, \quad i = 1, \dots, I_b + 2. \end{aligned} \quad (27)$$

In words, the estimate $(\hat{\lambda}, \hat{\alpha}_G)$ is chosen to make the model-generated conditional distribution of \tilde{K} match its empirical analog as closely possible.²⁴ The constraints on the

²³The boundary-corrected kernel function we use follows Karunamuni and Zhang [2008]. See Hickman and Hubbard [2015] for an in-depth discussion of its advantages and uses in structural auctions models.

²⁴An analogous estimator in the absence of the generalized Poisson assumption would be possible, but with additional complications. In the case where $\lambda = \{\lambda_0, \lambda_1, \lambda_2, \dots\}$ and $\lambda_k = \Pr[K = k]$, the main challenge is that only finitely many elements of λ can be estimated with finite sample size L . Thus, one

empirical objective function enforce boundary conditions and monotonicity of our parameterization for \hat{G}_B .²⁵

Finally, we separately estimate \hat{G}_R by a simpler method of moments procedure as

$$\begin{aligned} \hat{\alpha}_r &= \arg \min_{\alpha_r \in \mathbb{R}^{I_r+3}} \sum_{l=1}^L \left\{ [\hat{G}_R(r_l; \alpha_r) - \ddot{G}_R(r_l)]^2 \right\} \\ &\text{subject to} \\ \alpha_{r1} &= \ddot{G}_R(r), \quad \alpha_{r, I_r+3} = 1, \\ \alpha_{ri} &\leq \alpha_{r, i+1}, \quad i = 1, \dots, I_r + 2, \end{aligned} \tag{28}$$

where $\ddot{G}_R(r) = \sum_{l=1}^L \mathbb{1}(r_l \leq r) / L$ is the empirical CDF of reserve prices.

3.2.2. *Stage II*: Having these estimates in hand, we are able to directly re-construct the remaining structural primitives. Some Stage II objects will depend on the time discount factor, and where this is the case we so note by including δ as a parameter argument for the relevant functional.

could choose an upper bound $\bar{K}_L < \infty$, restrict $\lambda_k = 0$ for each $k > \bar{K}_L$, and define the following:

$$\begin{aligned} \{(\lambda_0, \dots, \lambda_{\bar{K}_L}), \alpha_b\} &= \arg \min_{(\lambda, \alpha_b) \in \mathbb{R}^{\bar{K}_L + I_b + 3}} \sum_{l=1}^L \left\{ [\tilde{\pi}(\tilde{k}_l | r_l; \lambda, \alpha_b) - \hat{\pi}(\tilde{k}_l | r_l)]^2 + [H(y_l; \lambda, \alpha_b) - \hat{H}(y_l)]^2 \right\} \\ &\text{subject to} \\ \sum_{k=0}^{\bar{K}_L} \lambda_k &= 1, \\ \alpha_{b,1} &= 0, \quad \alpha_{b, I_b+3} = 1, \\ \alpha_{b,i} &\leq \alpha_{b, i+1}, \quad i = 1, \dots, I_b + 2. \end{aligned}$$

The optimal choice of \bar{K}_L in finite samples involves a bias-variance tradeoff. The larger is \bar{K}_L , the less bias is introduced from restricting values of high-order elements of λ , but on the other hand the variance of the estimator will eventually increase with the ratio (\bar{K}_L / L) as well. A fully nonparametric estimator must also specify the rate at which \bar{K}_L should grow with the sample size. While interesting, the answers to these questions are beyond the scope of the current exercise, so we do not address them here. In a simpler setting than ours—a static bidding model of eBay laptop computer auctions with no binding reserve prices—Hickman et al. [2016] found strong evidence that the generalized Poisson assumption produced estimates of the bidder arrival process that could not be improved upon by relaxations of its parametric form, given their sample size of roughly 750 auctions.

²⁵One of the numerical benefits of using B-splines is their ease of incorporating shape restrictions, many of which can be imposed as simple linear constraints on the parameter values themselves. For example, under the Cox-de Boor recursion formula (with concurrent boundary knots), the only basis functions to attain a non-zero value at the boundaries are $\mathcal{F}_{b1}(\cdot)$ and $\mathcal{F}_{b, P_b+3}(\cdot)$, which both equal one at the upper and lower endpoints, respectively. Therefore, enforcing boundary conditions is equivalent to setting the first and/or last parameter value equal to the known boundary value(s) of the B-spline function, which also cuts down on computational cost by reducing the number of free parameters. Monotonicity is also quite simple: [de Boor, 2001, p.115] showed that a B-spline function $\hat{G}_B(b; \alpha_b)$ will be monotone increasing (decreasing) if and only if the parameters themselves are ordered monotonically increasing (decreasing). This avoids the necessity of imposing a set of complicated, nonlinear (and potentially non-convex) constraints on the objective function values, as would be the case with global polynomials, in order to enforce appropriate shape restrictions which ensure our solution is a valid CDF.

Before moving on, a word on spot market mechanisms is in order. Empirical work has often assumed that eBay employs a standard second-price auction mechanism. Recent work has shown that non-trivial differences exist due to bid increments, which we denote by $\Delta > 0$. As bids are received by the online server, typically the price is set equal to the second highest bid plus an increment, or $Y + \Delta$, similarly as in a second-price rule. However, a complication arises when the top two bids are within Δ of each other: in this case the second-price rule will not do, since the high bid represents the winner's maximal commitment to pay, and $Y + \Delta$ would exceed this amount. In that case, the price is set equal to the high bidder's own bid as in a first-price mechanism. Thus, eBay's pricing rule follows $p(b) = \min\{B_M + \Delta, b\}$.

Hickman [2010] proved existence and uniqueness of a monotone Bayes-Nash bidding equilibrium under this pricing rule in a static, one-shot auction where the number of bidders is known. This equilibrium involves demand shading because there is a positive probability that the winner's own bid will determine the price she pays. Hickman et al. [2016] showed, in a static bidding game with stochastic participation and no binding reserve prices, that a bidder's private value is identified from the distribution of bids through the equation:

$$v = b + \frac{G_{B_M}(b) - G_{B_M}[\tau(b)]}{g_{B_M}(b)}, \quad \tau(b) = \begin{cases} \underline{b} & \text{if } b \leq \underline{b} + \Delta \\ b - \Delta & \text{otherwise} \end{cases} \quad (29)$$

where $\tau(b)$ is a threshold function determining the point below one's own bid which, if the random variable B_M surpasses it, will trigger a first-price outcome.

Proposition 3.3 enables us to adapt equation (29) above for the static inverse bid function $\tilde{\beta}^{-1}$ in our model, but two adjustments are required since bidders in our spot-market game are best responding to the random variable Z rather than just to B_M . First, the boundary condition for a bidder's static decision problem is now $\tilde{\beta}^{-1}(\underline{b}) = \underline{b} + \frac{G_R(\underline{b}) - G_R[\tau(\underline{b})]}{g_R(\underline{b})}$, since the only way for bid \underline{b} to win is the event where $M = 0$. Second, letting $G_Z(z)$ denote the CDF of Z , our inverse static bid function is given by:

$$\hat{\beta}^{-1}(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) = \hat{v}_b = b + \frac{\hat{G}_Z(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) - \hat{G}_Z[\tau(b); \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r]}{\hat{g}_Z(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r)} \quad (30)$$

Using Stage I estimates we can construct the allocation rule and the distribution of Z :

$$\begin{aligned} \chi(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) &= \hat{G}_R(b; \hat{\alpha}_r) \sum_{m=0}^{\infty} \pi_M(m; \hat{\lambda}) \hat{G}_B(b; \hat{\alpha}_b)^m \\ &= \hat{G}_Z(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r), \quad b \geq 0 \end{aligned} \quad (31)$$

Equation (31) is a straightforward adaptation of (11) above, and note that we extend the domain of the function so that the right-hand side of the first line can also represent the distribution of the random variable Z . Taking into account the form of the hybrid pricing rule, we can also construct the payment function:

$$\begin{aligned} \rho(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) &= \underline{r} G_R(\underline{r}; \hat{\alpha}_r) + \int_{\underline{r}}^{\tau(b)} (t + \Delta) \hat{g}_Z(t; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) dt \\ &+ b \left(\hat{G}_Z[b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r] - \hat{G}_Z[\tau(b); \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r] \right). \end{aligned} \quad (32)$$

The first term on the right-hand side is for the event where a second-price rule is triggered, and the second is for the event where a first-price rule is triggered. Recall that we allow for the possibility that G_R has a mass point at the lower bound of its support.

Using the zero surplus condition, we can recover the per-period entry cost as:

$$\hat{\kappa} = \chi(\hat{v}_b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) \hat{v}_b - \rho(\hat{v}_b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) \quad (33)$$

as well as the dynamic inverse bid function and value function which are:

$$\hat{v} = \hat{\beta}^{-1} \left(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta \right) = \hat{v}_v \frac{1 - \delta \left[1 - \chi(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) \right]}{1 - \delta} - \frac{\delta \left(\rho(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r) + \hat{\kappa}_B \right)}{1 - \delta} \quad (34)$$

$$\hat{V} \left(v; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta \right) = \frac{\hat{v} - \hat{v}_b}{\delta} \quad (35)$$

The private value distribution is a best-fit B-spline function. We begin by specifying a grid of $J = I_v + 1$ points spanning the bid support, $\mathbf{b}_J = \{b_1, \dots, b_J\}$, and a knot vector \mathbf{n}_v that spans $[\hat{v}, \hat{v}] = \left[\hat{\beta}^{-1} \left(\hat{b}; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta \right), \hat{\beta}^{-1} \left(\hat{b}; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta \right) \right]$. This in turn defines our basis functions $\mathcal{F}_{v,i} : [\hat{v}, \hat{v}] \rightarrow \mathbb{R}$, $i = 1, \dots, I_v + 3$, from which we can now compute α_v :

$$\hat{\alpha}_v = \arg \min_{\alpha_v \in \mathbb{R}^{I_v+3}} \sum_{j=1}^J \left\{ \left[\hat{G}_B(b_j; \hat{\alpha}_b) - \sum_{i=1}^{I_v+3} \alpha_{vi} \mathcal{F}_{vi} \left[\hat{\beta}^{-1} \left(b_j; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta \right) \right] \right]^2 \right\} \quad (36)$$

subject to

$$\alpha_{v,1} = 0, \quad \alpha_{v,I_v+3} = 1,$$

$$\alpha_{v,i} \leq \alpha_{v,i+1}, \quad i = 1, \dots, I_v + 2.$$

Finally, the steady-state measure and distribution of new agents flowing into the market each period are:

$$\hat{\mu} = \left[1 - \pi \left(0; \hat{\lambda} \right) \right] G_R(\underline{b}; \hat{\alpha}_r) + \int_{\underline{b}}^{\bar{b}} g_R(r; \hat{\alpha}_r) \left(\sum_{k=1}^{\infty} \pi(k; \hat{\lambda}) \left[1 - G_B(r; \hat{\alpha}_b)^k \right] \right) dr \quad (37)$$

$$t_V(v; \hat{\lambda}, \hat{\alpha}_b; \hat{\alpha}_r, \delta) = \frac{\chi \left[\beta \left(v; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta \right); \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r \right] f_V(v; \alpha_v, \delta) \frac{\lambda_1}{1-\lambda_2}}{\hat{\mu}} \quad (38)$$

3.2.3. *Asymptotics and Standard Errors.* In the online supplemental appendix, we argue that our Stage I estimators $\hat{\lambda}$, $\hat{\alpha}_b$, and $\hat{\alpha}_r$ fall within the class of Generalized Method of Moments estimators. As such, it follows that they are consistent and asymptotically jointly normal with known formulae for computing standard errors. Since Stage II empirical objects are all smooth functions of Stage I parameters, it follows that they are also asymptotically normal, and their standard errors can be computed via the delta method. See Appendix B for a detailed discussion on computation of standard errors.

4. DATA AND RESULTS

We use a unique dataset on Amazon Kindle Fire tablet devices that we scraped from eBay during March through July 2013. Our scraping algorithm allowed us to capture all item listings on eBay during that period, and for each one we downloaded and stored various .html files including the item listing page and the bid history page. During the sample period we observed a total of 1,732 Kindle Fires listed as “new” (i.e., unused in a factory sealed box) or “new other” (i.e., unused in an unsealed box) for an average of 11.25 per day.

Each Kindle tablet had eight gigabytes of internal storage and a seven-inch screen with standard-definition resolution of 1024x600. The Kindle Fire tablets came pre-loaded with Amazon’s proprietary version of the Android-based operating system that prevents the user from accessing the full Android app market.²⁶ This makes the Kindle Fire a poor substitute for a standard tablet (e.g., Samsung Galaxy or Apple iPad) that can serve a dual role as a productivity tool or as a highly versatile consumer electronic device. Rather, the Kindle Fire is specifically designed to be a consumer access point exclusively to Amazon.com’s electronic media market, which includes e-books, periodicals, audio-books, music, and movies.²⁷ All transactions during the sample period were covered by the eBay Money Back Guarantee to insure consumers against potential unscrupulous sellers.²⁸

²⁶It requires specialized knowledge to uninstall the proprietary operating system, and doing so is costly since it invalidates all product guarantees issued by Amazon.com.

²⁷Amazon.com also maintains its own limited app market—primarily dedicated to entertainment and online shopping, but in June 2013 it contained less than one tenth the number of apps available in Apple’s App Store for iPhones or Google Play for Android devices. See [https://en.wikipedia.org/wiki/App_Store_\(iOS\)](https://en.wikipedia.org/wiki/App_Store_(iOS)); https://en.wikipedia.org/wiki/Google_Play; and https://en.wikipedia.org/wiki/Amazon_Appstore; information retrieved on 7/15/2016.

²⁸As of 7/15/2016, details on eBay’s consumer protection program were available at <http://pages.ebay.com/ebay-money-back-guarantee/questions.html>.

In order to further probe the homogeneity of our Kindle auctions sample, we manually examined 132 listings (a 10% sample of our downloaded raw .html files) from our final dataset. Unlike many other tablet devices, accessories are only rarely coupled with Kindle Fires: of these listings, only one mentioned an accessory (a Kindle case) that the seller had bundled into the sale. The vast majority of the listings with a condition of “New,” which eBay defines as factory-sealed in the box, had been opened. A common explanation was that the seller was merely checking that all of the parts (e.g., charging cord) are present. Only five of the surveyed “New” items explicitly mentioned that they are sealed in the box. We conclude from this that the “New” listings are best interpreted as items that are like new and essentially unused.

Because those listings with a low closing price are so crucial for identifying the participation cost, κ , we manually scrutinized all of these items. Although we identify $\underline{v} = \$66$ as the minimal observed sale price, we examined all listings with a closing price of less than \$80. Of these we removed listings that (for example) were selling Kindle accessories (e.g., cases) rather than the actual device or were offering a Kindle running a user-modified version of the Android OS. These atypical listings were largely isolated to the lower tail of the price distribution and were completely removed from the sample prior to analysis.

One final concern is that there may be residual auction-specific variation which our manual survey missed, and which is not included in our econometric model. Unobserved heterogeneity (UH)—some auction characteristic that bidders see but the econometrician does not—is a common problem, and various approaches have been developed to deal with it (e.g., see seminal work by Krasnokutskaya [2011]). Each approach assumes bidder valuations are separable in the UH and the idiosyncratic component, which makes it possible to deconvolve UH from agent-specific variation in bids. More recently, Roberts [2013] proposed a method to correct for UH when only one bid is observed per auction. Since we can only be confident that the highest losing bid in each of our auctions is fully reflective of equilibrium strategies (see discussion below), Roberts [2013] is the most relevant paper to the current context. Under the assumption that seller reserve prices are also a separable function of the UH variable, he shows that one can use joint movement in reserve prices and bids to deconvolve the UH and identify private valuations.

In our data we observe non-trivial variation in sellers’ reserve prices with roughly one third of them being binding for a positive fraction of the bidder population. Therefore, one might reasonably suspect that if UH is present then higher values of the unobserved

characteristic prompt sellers to increase reserve prices.²⁹ A necessary condition for UH in the Roberts model is co-movement of bids and reserve prices, which is testable. We find in our data that the correlation between seller reserve price and the highest losing bid is very small in a practical sense, at -0.015, and that it is also statistically indistinguishable from zero. We interpret the combined information above—lack of correlation between bids and reserve prices, evidence from our manual survey of .html pages, uniform buyer insurance, proprietary operating system and limited app market, and uniform characteristics of the Kindle Fire tablets—as evidence consistent with our assumption of a homogeneous goods market with no close substitutes. These characteristics of the eBay Kindle data allow us to avoid significant complications covered by other work, such as identifying UH or complex substitution patterns (see Backus and Lewis [2016]), and instead focus on questions of bidding behavior, allocative efficiency, and market design.

4.1. Practical Concerns.

4.1.1. *Intra-Auction Dynamics.* For each auction listing, we observe the timing and amount of each bid submission as well as the bidder identity that goes with the bid. As previous empirical work has recognized, one challenge for interpreting eBay data is a large number of implausibly low bids early on in the typical auction. Many bidders place repeated bids, often within a few dollars or cents of each other, and then become inactive long before the posted price approaches a reasonable level. Some bidders may engage in non-equilibrium cheap-talk before bidding based on best-response calculations or participate flippantly to pass time while web surfing. Empirically, a significant fraction of observed bid amounts, especially those submitted early in the life of the auction, fall too far below realistic transaction prices to be taken seriously. The question of intra-auction dynamics is broad, complicated, and beyond the scope of this work.³⁰ In our case, inter-auction dynamics are the primary concern for answering our research questions on allocative efficiency and market design.

To deal with observed early low bids, we adopt the approach of Bajari and Hortaçsu [2003] by partitioning individual auctions into two stages. During the first phase bidders may submit cheap-talk bids that are viewed as uninformative of the other bidders' final bids and the final sale price. The second stage is treated as a sealed-bid auction as

²⁹Note that the model of Roberts [2013] does not require a specific theory to rationalize sellers' choice of reserve prices. Rather, it assumes only that reserves are a monotone separable function of UH. For example, it doesn't matter whether reserves are chosen to optimize projected revenues or whether they are chosen to hedge against the risk of selling at an unacceptably low price, since both scenarios would satisfy monotonicity. Our structural estimates suggest the latter model as most plausible.

³⁰The leading attempts in the literature to formalize intra-auction dynamics are Nekipelov [2007] and Hopenhayn and Saeedi [2016].

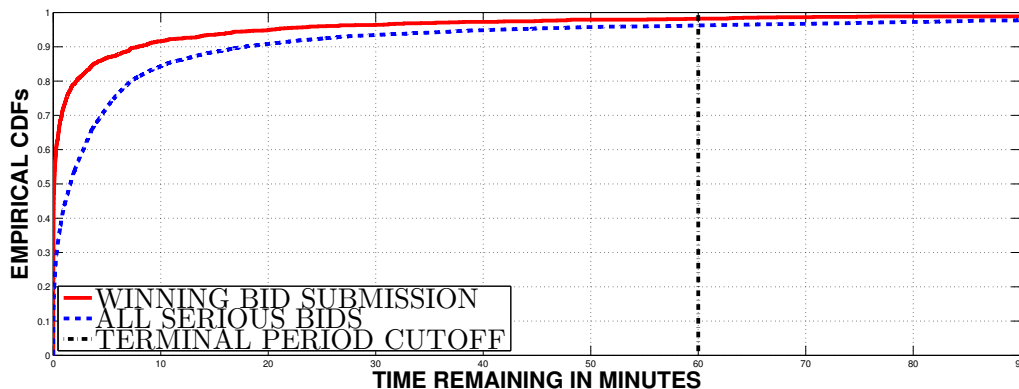


FIGURE 3. Empirical Distributions: Time Remaining when Bids are Submitted

per our model of Sections 2 and 3.³¹ Finally, consistent with the previous section, the ordering of bidders' submission times is assumed to be random rather than coordinated.

This requires us to take a stand on differentiating between bids that are a meaningful part of competition and those that are superfluous. We define a *serious bid* as one that affects the price path within the second stage of an auction. Likewise, a *serious bidder* is one who is observed to submit at least one serious bid. Of course, the possibility always exists that some bidders who are determined to be non-serious by the above criterion had serious intent to compete for the item, but were priced out before submitting their planned, serious bid during the terminal stage. This is, however, part of the problem that our model of the adjusted filter process solves (i.e., observed participation by serious bidders is a lower bound on actual participation). Finally, note that our definition of serious bidding will also count the top two submissions from within the first stage of the auction as these bids fix the price at the start of the second stage of the auction. This allows us to avoid drawing too sharp a distinction between the cheap-talk stage and the terminal stage of the auction, since some serious bidders' submission times may still occur early in the life of the auction.

We specify the terminal period as the last 60 minutes of an auction, during which we see an average of 4.01 observed serious bidders per auction. Figure 3 shows the empirical distribution for time remaining when the winning bid was submitted, which occurs within the final 60 minutes in over 95% of auctions in the sample. The figure also shows the empirical distribution for time remaining across all serious bid submissions in the sample. These figures are not sensitive to alternate specifications of the terminal

³¹While eBay auctions that run for several days can attract bids prior to the final moments, the vast majority of eBay auctions are won by bidders who bid in the final moments and the terminal behavior of the price path is largely independent of overall auction duration. This phenomenon was first documented empirically by Roth and Ockenfels [2002].

TABLE 1. Descriptive Statistics

Variable	Mean	Median	St. Dev.	Min	Max	# Obs
Time Remaining (minutes)						
Winning Bid Submission:	6.69	0.11	38.31	0.00	593.30	1,460
High Loser Bid Submission:	12.49	0.56	52.85	0.00	604.35	1,397
Observed Participation						
\tilde{N} (serious bidders only):	4.01	4	1.82	0	12	1,462
Monetary Outcomes						
Sale Price:	\$124.96	\$125.00	\$17.74	\$67.00	\$190.00	1,460
Highest Losing Bid:	\$123.84	\$124.50	\$17.34	\$66.00	\$189.50	1,397
Seller Reserve Price:	\$33.56	\$0.99	\$45.27	\$0.99	\$175.00	1,462

period cutoff. If it is chosen as 80 minutes the mean number of serious bidders becomes 4.25, and if it is chosen as 40 minutes the mean number of serious bidders becomes 3.67.

Given our algorithm for distinguishing between serious and non-serious bid submissions, there remains one final challenge. Bidders may choose to submit their strategic bid at once to the server and make use of eBay’s automated proxy bidding, or they may choose to incrementally raise their bid submissions up to the level of their strategic bid on their own. Roughly one third of serious bidders are observed to engage in incremental bidding. Since it is unclear how to interpret each individual bid submission that affects the terminal price path, we assume that only the highest losing bid is fully reflective of equilibrium play. This leaves us with the three data points from each auction that we need for identification: \tilde{k}_l , the observed number of serious bidders; r_l , the seller’s reserve price; and y_l , the highest loser bid from auctions with at least two bids. After dropping .html pages for which our software was unable to extract data because of formatting problems, we have 1,462 total auctions, 2 of which logged no bids, and 1,397 of which had 2 or more observed bidders so that we observed a highest losing bid. Table 1 displays descriptive statistics on bid timing, observed participation, sale prices, and highest losing bids.

4.1.2. *Model Tuning Parameters.* Before implementing the estimator there remain several free parameters from the previous section to pin down. The most important of these are the knot vectors \mathbf{n}_b , \mathbf{n}_r , and \mathbf{n}_v . We adopt the convention that knots will be uniformly spaced, which then reduces the problem to choosing values for I_r , I_v , and I_b that dictate the number of knots to use in the relevant B-spline function. For the first two we first

choose a grid of uniform points in $[0, 1]$ (quantile rank space), and then we map these back into R space (or V space) using the empirical quantile functions. This procedure ensures that the influence of the data is spread evenly among the various basis functions. For \mathbf{n}_b , we chose knots that are uniform in bid space. The reason for this is that α_b directly parameterizes the parent distribution \hat{G}_B , but in our estimator we are matching the empirical moments of the order statistic distribution H without knowing the quantiles of G_B ex ante.

In Stage I we chose $I_b = 10$, and we partitioned the reserve price support by the quintiles of the empirical conditional distribution $\check{G}_R(r|R > \underline{r})$, meaning $I_r = 5$.³² This gives us a total of 13 parameters for \hat{G}_B and 8 for \hat{G}_R . We chose $I_v = 15$ knots at the quantiles of the distribution $\hat{G}_B \circ \hat{\beta}$, which is known from Stage I. We chose $I_v > I_b$ because \hat{F}_V must conform to the nuances induced by all first-stage parameters in order to accurately represent the implied private value distribution. We find that these choices provide a good fit to the data and that adding more parameters renders little benefit.³³ The interested reader is directed to Figure 9 in the online appendix, which displays the complete set of knots and B-spline basis functions that make up \hat{G}_B , \hat{G}_R , and \hat{F}_V . This figure is also meant to give the reader a sense for how knot location choice alters the form of the basis functions.

The final free parameter is the time discount factor, δ . As in many other empirical contexts, this part poses a difficult challenge. Luckily, δ does not enter Stage I estimation, so all of the necessary building blocks to compute the final structural primitives will be unaffected. Several Stage II objects are also unaffected, including the win probability $\chi(\cdot)$, the expected payment function $\rho(\cdot)$, the per-period bidding cost $\hat{\kappa}$, and the exogenous, per-period measure of new agents flowing into the market $\hat{\mu}$. However, the remaining objects depend on δ . The objects affected by δ include the dynamic bid function $\beta(\cdot)$, the value function $\mathcal{V}(\cdot)$, and the steady-state private value distributions for market participants $F_V(\cdot)$ and new entrants $T_V(\cdot)$. There is an intuitive reason why: these objects tell us something about the opportunity cost of losing today, and δ plays a pivotal role in shaping this opportunity cost by determining agents' attitude toward present versus future consumption.

In lieu of taking a stand on the particular value of δ applicable to our study, we present results both here and in our counterfactual setting for a range of values of δ . Where possible, we provide statistics that are stable across choices of δ . For example, instead

³²The conditioning is due to the mass point at the lower bound.

³³A fully semi-nonparametric estimation routine based on B-splines would involve specifying a rule for optimal choice of I within finite samples and the rate at which I should increase as the sample size $L \rightarrow \infty$. This is an interesting econometric question, but one which is beyond the scope of this paper.

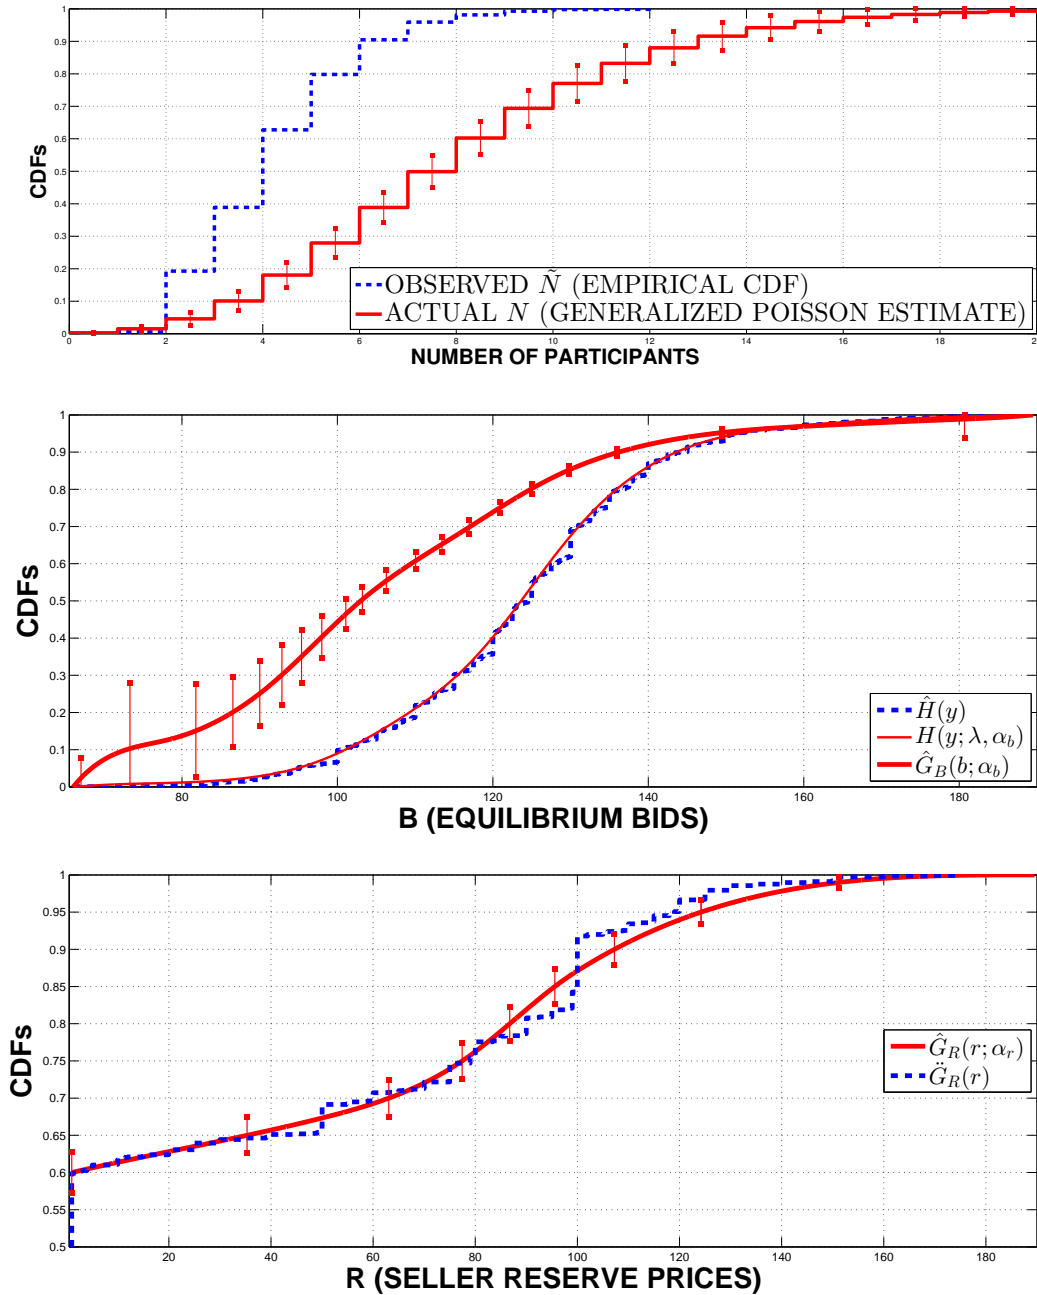


FIGURE 4. Stage I Estimates

of providing a dollar value for deadweight loss, which is sensitive to δ , we present deadweight loss as a percentage of the buyer’s value, which is stable across different choices of δ .

4.2. **Estimates.** Table 2 displays point estimates and standard errors for readily interpretable parameters, including the market tightness parameters, the per-period bidding cost, and the per-period measure of new entering agents. Figure 4 depicts point estimates

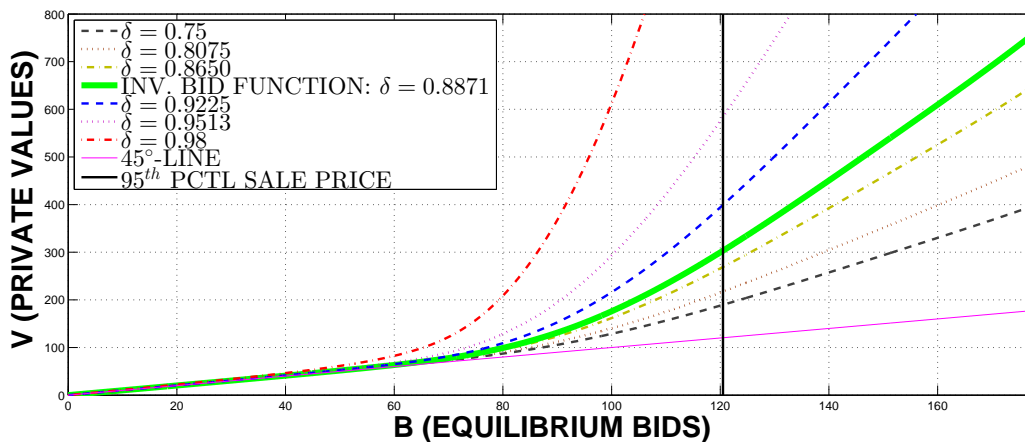


FIGURE 5. Inverse Bid Function Estimates Given Various Values of δ

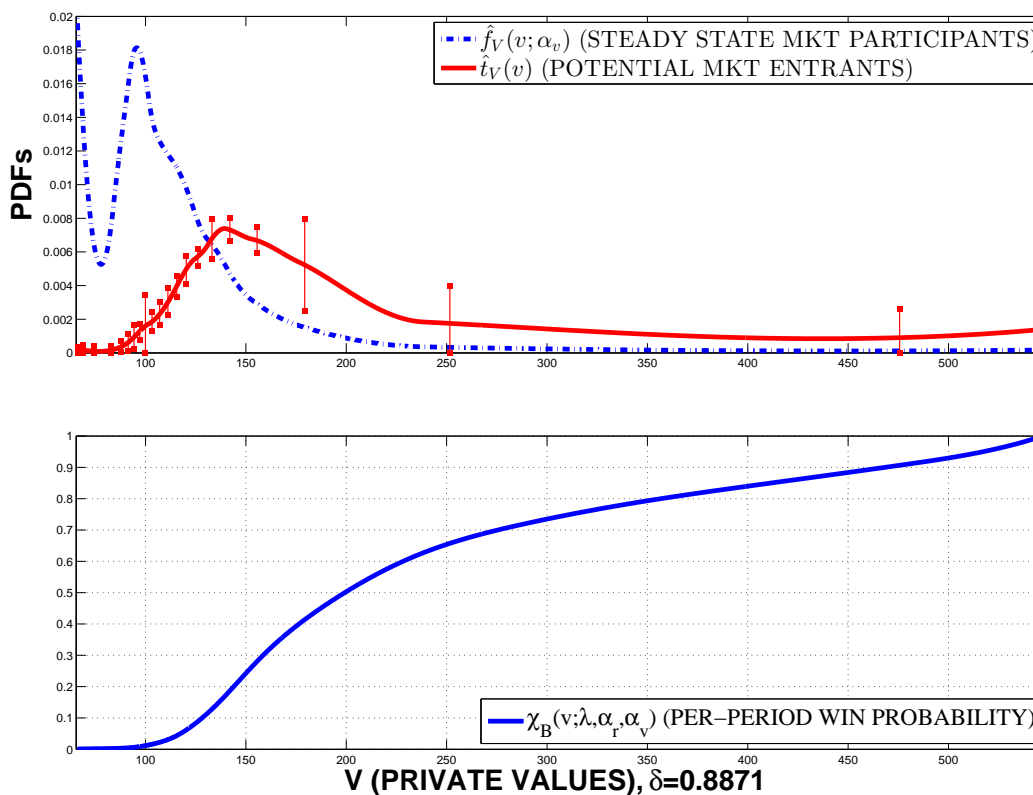


FIGURE 6

for our Stage I distribution estimators (thick, solid lines), point-wise confidence bounds for a selected grid of domain points (vertical box plots), and empirical distributions being matched by the model (thick, dashed lines). The first panel shows the empirical CDF of observed bidders \tilde{K} and the estimated distribution of total auction-level participation K . As the figure demonstrates, failing to account for unobserved bidders within the spot

TABLE 2. Estimation Results

Variable:	λ_1	λ_2	κ	μ
Point Estimate:	5.9100	0.2579	0.0654	0.9649
Standard Error:	(0.384)	(0.058)	(0.0174)	(0.0261)

market sample selection process would lead to a very different view of the distribution of auction participation. This substantial difference shows up in both the mean—4.07 for observed bidders per auction versus 7.96 for actual bidders per auction—and also in the variance—3.19 for observed bidders and 14.46 for actual bidders. The lower two panels provide an idea of the model fit. The middle one depicts model fit for the distribution of the highest loser bid and includes an extra plot for the model-driven $H(y; \hat{\lambda}, \hat{\alpha}_b)$ distribution, which is derived from both the market tightness and parent bid distribution parameters. The lower panel depicts model fit for the seller reserve price distribution. Note that in both cases, the B-spline functions provide a very good fit to the underlying data. The difference between the two cases is that in the latter our B-splines parameterize the distribution \hat{G}_R , which is directly matched to its empirical quantiles, whereas in the former, we parameterize \hat{G}_B and then indirectly match the moments of the implied order statistic distribution H .

Figure 5 presents the dynamic inverse bid functions $\hat{\beta}^{-1}(b; \hat{\lambda}, \hat{\alpha}_b, \hat{\alpha}_r, \delta)$ which we estimate for a uniform grid of values of the time discount factor δ between 0.75 and 0.98. We also include an additional value at 0.8871 taken from an experimental study by Augenblick, Niederle, and Sprenger [2016] where they elicited hyperbolic time discount parameters at the daily level from college students.³⁴ Recall from Figure 1 that the vast majority of demand shading is driven by the option value of returning to the market in future periods if one does not win today. This continuation value is primarily driven by three things: the equilibrium bid distribution \hat{G}_B , the market tightness parameters λ , and the discount factor δ . Figure 5 depicts the important role of this third piece. Since δ determines bidders' attitudes toward trading off today's consumption for tomorrow's, a greater degree of patience requires larger values of v to rationalize observed bids. Recalling that δ is a daily discount factor, if we adopt a value of 0.98 then the 95th percentile of the private value distribution is over \$1,300, which we consider to be implausibly high.

³⁴Another related study by Burks, Carpenter, Götte, and Rustichini [2012] elicited daily time discounting preferences from professional truck drivers in a field experiment using real monetary incentives distributed through their employer. Burks et al.'s data led to an estimated average daily discount factor of 0.8921.

TABLE 3. Mean Private Values and Information Rents For Various δ

Discount Factor δ :	0.75	0.81	0.87	0.8871	0.93	0.95	0.98
Mean Private Value:	\$48.57	\$51.29	\$56.32	\$59.63	\$68.83	\$86.15	\$153.24
Mean Winner Private Value:	\$208.39	\$230.98	\$269.26	\$293.58	\$358.55	\$474.05	\$875.56
Mean Winner Information Rent:	\$54.66	\$69.11	\$94.08	\$111.09	\$157.44	\$245.84	\$583.91
Mean Information Rent Percentage:	26.23%	29.92%	34.94%	37.84%	43.91%	51.86%	66.69%

Table 3 displays various descriptive statistics derived from Stage II estimates, including average private values, average private values of winners, and information rents (i.e., the difference between the winner’s private value and the spot-market price). The last row of the table shows information rents as a fraction of the winner’s private value, on average. Finally, Figure 6 presents other Stage II estimates related to the distribution of buyer values. The upper pane displays the PDF of the distribution of market participants’ private values in steady state under our preferred specification, $\hat{f}_V(v; \alpha_v, \delta = 0.8871)$ (dash-dot line), as well as the type distribution for new market entrants each period, $\hat{t}_V(v; \hat{\lambda}, \hat{\alpha}_b; \hat{\alpha}_r, \delta = 0.8871)$ (solid line), with point-wise confidence bounds (vertical box plots). The PDFs t_V and f_V are tied together by the win probability, χ , depicted as a function of v for comparison. Although there are many buyers in the market with low values in steady state, our model suggests that relatively few of these agents enter the market each period. However, those low-value buyers that do enter must stay in the market for a long period of time before winning, as indicated by the function χ . The long delay between entry and purchase means these low value buyers accumulate in the market out of proportion to their presence in the distribution of new entrants. On a related note, the delay between entry and trade implies that the lower-value buyers are the ones most significantly affected by the per-period participation cost, $\hat{\kappa} = \$0.065$.

In comparing estimates of f_V and t_V in Figure 6, two important differences should be noted. First, since f_V depicts the type distribution for all market participants—including players remaining from previous periods after failing to win a spot-market auction—it represents a measure $\lambda_1/(1 - \lambda_2) = 7.96$ of agents (recall that sellers are assumed to have measure 1). On the other hand, t_V describes the type distribution of the measure

$\mu = 0.9649$ of new agents that enter the market each period in order to maintain the steady state. Another way of interpreting μ is that each spot-market auction has roughly 3.5% probability of resulting in no sale, either because $k = 0$ players were matched to it or (much more likely) because the reserve price was too high.

The second important difference between f_V and t_V has to do with the probability that each agent type v will transact and exit the market. Under f_V there is a relatively large mass of low-value bidders, who are not very likely to win each period, and so in turn they tend to pile up in the market and remain for many periods until finally winning an auction. On the other hand, t_V depicts a selected set of buyers who move in and out of the market at much higher frequency, on average, because they have higher private values, and are much more likely to win in the spot market in a given period.

5. COUNTERFACTUALS

We now perform three counterfactual analyses to investigate the economic implications of our structural model. The first explores market efficiency. The second decomposes the relative importance of what we refer to as *platform composition* (PC) effects (i.e., market entry/exit when market conditions change) and *dynamic incentive* (DI) effects (i.e., when bidding behavior changes in response to shifts in opportunity costs). The third counterfactual exercise investigates optimal starting price choice in order to assess whether there are significant costs to setting the starting price at the (suboptimal) minimal value of \$0.99. For notational simplicity we omit the parameter arguments of primitive structural functionals unless needed for clarity.

Before proceeding, we would like to briefly describe the algorithm used to compute the counterfactuals. For expositional clarity, we focus on the SPA pricing rule. The structural primitives of our model are μ , t_V , κ , G_R , and λ_2 .³⁵ These structural parameters remain fixed when performing our counterfactual exercises unless otherwise noted. As discussed in Section 2.2, the following conditions pin down \underline{v} and β and aggregate variables C , λ_1 , and F_V for a SPA spot market mechanism. Letting \bar{v} denote the highest

³⁵We would have liked to allow λ_2 to be endogenized, but we did not see any obvious economic structure that would naturally pin this variable down endogenously.

value buyer in the data, we have the system:

$$\mu t_V(v) = \chi(\beta(v)) f_V(v) \mathcal{C} \quad (39)$$

$$\beta(v) = v - \delta \mathcal{V}(v) \quad (40)$$

$$\mathcal{V}(v) = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \int_{\underline{v}}^v (1 - \chi(\beta(s)))^{\tau-t} \chi(\beta(s)) ds \quad (41)$$

$$\chi(\beta(\underline{v})) \underline{v} - \rho(\beta(\underline{v})) = \kappa \quad (42)$$

$$F_V(\bar{v}) = 1 \quad (43)$$

Equation 39 requires that the distribution and measure of buyers exiting after winning an auction equals the distribution and measure of buyers flowing in. Equations 40 and 41 pin down the bidding strategy and the value function. Equation 40 will take the form of an ordinary differential equation in non-SPA spot markets, but this is a relatively trivial modification to our structure. Equation 42 requires that the lowest value buyer that enters be indifferent to entering, and Equation 43 requires that the steady-state distribution of types be properly normalized. The function $\chi(b)$, which appears in several of these equations, is defined by the allocation mechanisms. In our single unit auctions we have:

$$\chi(b) = G_R(b) \sum_{m=0}^{\infty} \pi_M(m) F_B(\beta^{-1}(b))^m$$

In practice, our software uses a bisection algorithm to search for the equilibrium value of \underline{v} . Given \underline{v} , we let the auction size parameter, λ_1 , adjust so that Equation 42 holds. Since $\mathcal{C} = E[K] = \frac{\lambda_1}{1-\lambda_2}$ by assumption, we immediately have \mathcal{C} . Given \underline{v} , λ , and \mathcal{C} we can solve Equations 39 - 41 to obtain candidate values for f_V , β , and \mathcal{V} . Finally, if Equation 43 fails to hold for the candidate f_V (i.e., if the steady-state distribution of types is not properly normalized), then we adjust our guess for \underline{v} and repeat the process.³⁶

5.1. Welfare Comparisons. Throughout this section we adopt the usual notion of auction efficiency as the tendency for goods to be allocated to those who value them most within a given period. Even when the spot-market mechanism is efficient within a given auction, dynamic auction platforms with search frictions still exhibit two related sources of inefficiency. First, there is the chance that a high-value buyer that ought to receive the good in an efficient allocation is competing against another high-value buyer, so one of them cannot receive the good. Second, an auction may fail to attract any high-value buyers, which means a low-value buyer will receive the good when she would not under

³⁶Since we do not have any data on the distribution of buyer values below the \underline{v} observed in the data, we can only perform counterfactuals that yield a value of \underline{v} that is weakly larger than that observed in the data.

an efficient outcome. The first case is one in which there is “too much” competition in the auction, while the second case is one in which there is “too little” competition.

5.1.1. *“Model Anemic” Inefficiency Calculations.* In this section we use only our Stage I estimates to bound the percentage of auctions resulting in an inefficient sale. We refer to these calculations as “model anemic” since they do not rely on our equilibrium bidding model and thereby employ the fewest possible assumptions. Our model-anemic calculations rely only on our filter process model to correct for sample selection in the observed number of bidders in each auction.

To proceed, we must first find the cutoff v_{eff} that separates high-value buyers that ought to receive the good in an efficient allocation from lower-value buyers that ought not. Since the buyer-seller ratio is $\lambda_1/(1 - \lambda_2)$, the efficient allocative cutoff in private value space is defined by $v_{eff} \equiv F_V^{-1}\left(1 - \frac{1-\lambda_2}{\lambda_1}\right)$. However, since quantile orderings are invariant to monotone transformations, we can re-define this cutoff in bid space (where the raw data live) as $b_{eff} \equiv G_B^{-1}\left(1 - \frac{1-\lambda_2}{\lambda_1}\right)$. Intuitively, if the highest losing bid in a given auction exceeds b_{eff} , then the corresponding bidder would receive the good in an efficient allocation. We find that 28.47% of the auctions in our sample satisfy this criterion. For each high-value bidder who loses an auction there is a low-value bidder in some other auction who inefficiently wins, so high-value buyers losing and low-value buyers winning are simply two sides of the same coin.³⁷

This measure is only a lower bound on the frequency of inefficiency because without observing more bids, we cannot account for auctions where two or more losing bids surpassed b_{eff} . Another disadvantage of the model-anemic approach is that it offers no way of measuring the magnitude of unrealized gains from trade. Such an undertaking requires one to quantify private values that underpin observed bids.

5.1.2. *Structural Welfare Calculations.* Our full Stage II structural estimates allow us to get a more complete idea of the frequency and magnitude of market inefficiency. First, using Equation 7 we can compute the precise frequency of inefficient allocations as the fraction of all transactions involving low-value bidders:

$$\Pr[v_{winner} < v_{eff}] = C \int_{\underline{v}}^{v_{eff}} \chi(\beta(s)) f_V(s) ds.$$

Note that this measure is invariant to choice of the time discount factor δ . Our point estimates imply that 35.89% of Kindle auctions on eBay end with an inefficient outcome.

³⁷There is also a very small fraction of auctions that result in no sale due to high reserve price or $K = 0$ by random chance, but these scenarios happen too infrequently to be a significant source of welfare loss, so we ignore them until the next section where our measurements use the full structural model.

Deadweight loss calculations in levels will be sensitive to the choice of δ . In order to address this problem, we adopt the following measure, which we refer to as the *efficiency ratio*:

$$\mathcal{E}_{u,\delta} = \frac{\mathcal{C} \int_{\bar{v}} s \chi_u(\beta_u(s)) f_{V,u}(s) ds}{\mathcal{C} \int_{\bar{v}_{eff}} s f_{V,u}(s) ds}$$

The numerator is the realized gains from trade in our market (within a given period), and the denominator represents gains from trade generated by a fully efficient allocation. The u subscript denotes number of units involved in each auction listing for our counterfactual centralization analysis below; for now we fix $u = 1$. By expressing surplus as a fraction of total possible surplus, the separate influences of δ in the numerator and denominator largely cancel out and we get a measure that is stable across different assumptions on time discounting (see alternative calculations displayed in the first row of Table 4). We also compute the efficiency ratio under a hypothetical lottery system, denoted $\mathcal{E}_{lott,\delta}$, as the minimum efficiency benchmark (see the last row of Table 4).

With these definitions in hand, our point estimates imply that the fraction of total deadweight loss is $1 - \mathcal{E}_{1,0.8871} = 0.135$ under our preferred specification. To put this number into context, deadweight loss under a lottery system is estimated to be $1 - \mathcal{E}_{lott,0.8871} = 0.53$, meaning that eBay's auction market platform achieves only 76% of total gains from trade above the lottery benchmark. Note, however, that this is only a "partial equilibrium" assessment; were a social planner with complete knowledge of the bidder values to implement the efficient allocation each period, then the steady-state distribution of buyers' values and the buyer-seller ratio would change. However, we believe our figures have the benefit of giving a sense of the welfare losses while imposing minimal structural assumptions on the estimates.

5.1.3. Counterfactual Market Centralization. We now consider the extent to which inefficiencies can be mitigated by changing the market structure to one in which the same number of Kindles are allocated each period, but using fewer u -unit, uniform-price auctions with $u \geq 2$. Since new Kindles are relatively homogenous products, we think it is reasonable to assume that buyers view them as nearly perfect substitutes for one another. This suggests that our proposal to take steps toward more efficient market centralization using multi-unit auctions is feasible. In product categories where the items are not perfect substitutes (e.g., used cars), the implications of selling disparate products in a multi-unit auction become much more difficult to formalize. However, our estimates provide a sense of the efficiency loss generated by search frictions when selling items through decentralized, single-unit auctions as opposed to more centralized market mechanisms.

Several aspects of our model need to be slightly adjusted in the multi-unit auction setting. We subscript the endogenous quantities in our counterfactual equilibrium with u, δ to denote the degree of centralization and the choice of time discount factor. First, each u -unit auction attracts a number of bidders K_u distributed as a generalized Poisson random variable with expected value

$$E[K_u] = \frac{\lambda_{1,u}}{1 - \lambda_2} = uC \quad (44)$$

We assume that λ_2 , the dispersion parameter, is fixed at the estimated value and allow $\lambda_{1,u}$, the size parameter, to adjust so that (44) is satisfied in our counterfactual equilibria.

In our status quo model, we assume that each seller draws an independent reserve price from G_R . In the multi-unit context, we assume that a single reserve price is drawn from G_R for each u -unit auction, and that reserve price applies to all u units being allocated in that auction.³⁸ Each bidder submits a bid to the auction to which she is matched, and the u highest bids that are larger than the auction's reserve price win an item. Each winning bidder then pays a sum equal to the largest of the $(u + 1)^{th}$ highest bid and the reserve price. The implied probability of winning in a u -unit auction is:

$$\chi_u(b) = G_R(b) \left[\sum_{m=0}^{u-1} \pi_M(m) + \sum_{m=u}^{\infty} \pi_M(m) \binom{m}{u-1} F_B(\beta^{-1}(b))^{m-u+1} \right] \quad (45)$$

which we use in conjunction with Equations 39 - 43 and Equation 44 to compute the SCE for a u -unit auction.

One of the general takeaways from our research is that understanding the impact of platform market design on participation decisions is crucial. The social planner's welfare calculus will be strongly influenced by changes in entry behavior (e.g., how many low-value buyers leave the market?) and the steady state-distribution of private values for market participants (e.g., how many low-value bidders accumulate in the market when they are less likely to win an item?) Our model allows us to handle these questions by computing the counterfactual, steady-state SCE when the platform uses u -unit spot-market mechanisms. We find that \underline{v} increases as the market becomes more centralized (i.e., as u grows) since player types with very low values will see a decrease in their probability of winning as market allocations become more efficient.

Table 4 provides results for counterfactual efficiency ratio statistics for $u \in \{1, 2, 4, 8\}$. Recall from above that the efficiency ratio compares gains from trade in a single period of a u -unit model with the welfare generated by the efficient allocation from clearing the market once per period with a single, large, multi-unit, uniform-price auction.

³⁸We include a reserve price to make comparisons between the status quo setting and the uniform price setting we study here.

TABLE 4. Counterfactual Efficiency Ratios $\mathcal{E}_{u,\delta}$

# Units Per Listing	Discount Factor $\delta =$						
	0.75	0.80	0.86	0.88	0.92	0.95	0.98
1	0.89	0.88	0.87	0.86	0.85	0.84	0.82
2	0.92	0.92	0.91	0.91	0.91	0.90	0.89
4	0.94	0.94	0.94	0.94	0.94	0.93	0.93
8	0.95	0.95	0.95	0.95	0.95	0.95	0.95
Lottery	0.58	0.54	0.49	0.47	0.41	0.35	0.26

We would like to draw attention to two features of our results. First, the efficiency ratios are remarkably stable across different specifications of the time discount factor δ . Second, the vast majority of possible gains from centralization can be realized by 2- or 4-unit uniform-price auctions, so there is little need to shift towards a fully centralized market.

One might naturally expect that if eBay could re-design their platform market to increase allocative efficiency, then it ought to be able to benefit by capturing some of the increased gains from trade.³⁹ However, a careful examination of the moving parts within the model indicates that the sign of the effect on revenue is ambiguous. On the one hand, a given bidder with a value above the new participation cutoff \underline{v}_u faces fewer competitors in the market. On the other hand, her remaining competitors also value the object more highly on average. This complex combination of effects make it difficult to derive a priori predictions on bidding behavior and the resulting effects of revenue. Table 5, which describes the mean revenues generated per-auction as a function of u , demonstrates that the average sale price actually falls as u increases.

To help explain why revenue drops as efficiency rises, Figure 7 plots the probability of winning for each type of agent and the equilibrium bid function for the $u = 1$ (solid line) and $u = 8$ (dashed line) market structures. The win probability plot reveals that for most agents (especially those most likely to win), increasing market efficiency raises the probability that they will win a spot-market auction within a given period. This raises their future continuation values and therefore the opportunity cost of winning an auction today. This in turn reduces their bids by promoting further demand shading as shown in the second panel of Figure 7. Reduced bids then translate into decreased revenues for both sellers and eBay, which currently charges the sellers a percentage commission

³⁹The literature on optimal auctions suggests that efficiency reducing reservation prices can increase revenue. It is not clear if/how this result applies to our counterfactual other than the general sense that allocative efficiency and revenue are in tension.

TABLE 5. Counterfactual Mean Auction Revenues

# Units Per Listing	Discount Factor $\delta =$						
	0.75	0.80	0.86	0.88	0.92	0.95	0.98
1	\$115.05	\$114.90	\$114.71	\$114.59	\$114.37	\$114.10	\$113.65
2	\$112.79	\$112.33	\$112.97	\$112.72	\$112.26	\$111.96	\$112.28
4	\$111.73	\$111.21	\$111.09	\$110.84	\$110.40	\$110.31	\$111.30
8	\$110.05	\$109.54	\$109.12	\$108.94	\$108.64	\$108.81	\$109.75

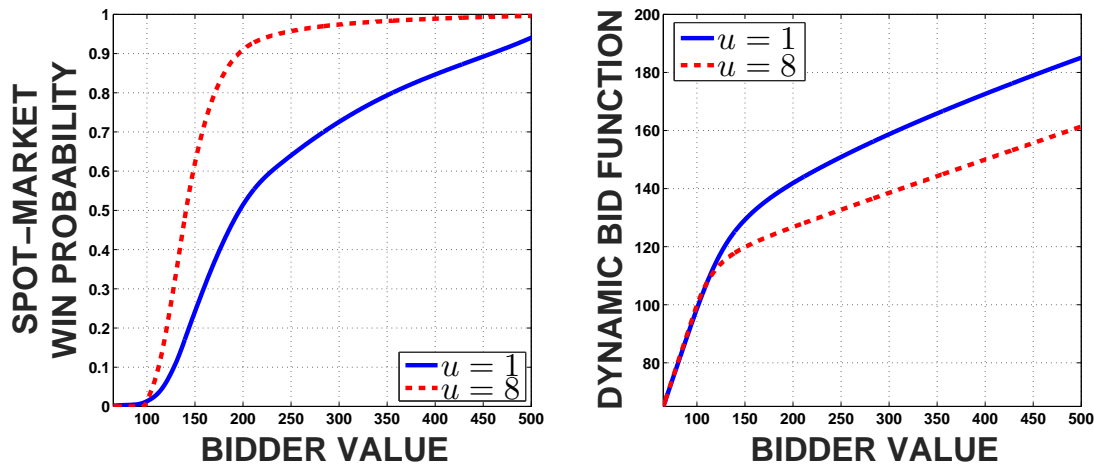


FIGURE 7. The Efficiency-Revenue Link

on auction revenue. This highlights an interesting point: what is good for bidders and market welfare is not necessarily good for platform market designers like eBay.

Finally, we would like to highlight the effect of centralization on the average lifetime participation costs (ALPC) paid by the agents, which are summarized in Table 6. For an agent of type v , the expected lifetime participation costs paid by the agent is equal to $\kappa/\chi(\beta(v))$. The ALPC refers to this quantity averaged over the steady-state distribution of buyers multiplied by the buyer to seller ratio:

$$\mathcal{ALPC}_{u,\delta} = C \int_0^\infty \frac{\kappa}{\chi(\beta(v))} f_{V,u,\delta}(v) dv$$

The ALPC is described in Table 6.

There are two effects at work. First, when markets centralize, the ALPC paid by a participant before winning an item goes slightly up on average. For example, when $\delta = 0.88$ and $u = 1$,⁴⁰ the participants paid on average \$8.98 each over their lifetimes in

⁴⁰Since the estimation of κ is independent of δ , the figures discussed below are essentially identical for all choices of δ .

TABLE 6. ALPC Effects of Centralization

# Units Per Listing	Discount Factor $\delta =$						
	0.75	0.80	0.86	0.88	0.92	0.95	0.98
1	\$66.57	\$67.77	\$68.95	\$69.58	\$71.27	\$72.97	\$75.92
2	\$47.80	\$50.03	\$52.33	\$54.30	\$58.08	\$60.26	\$61.33
4	\$34.59	\$36.09	\$37.33	\$38.07	\$40.14	\$44.86	\$61.51
8	\$22.55	\$23.84	\$26.05	\$27.73	\$32.46	\$41.49	\$65.59

the market. When $\delta = 0.88$ and $u = 8$, the participants paid on average \$9.50 over the course of their participation in the market. The average participation cost per bidder is pushed up by the fact that lower value agents must wait even longer (on average) before winning an item and exiting the market.

The larger effect is that fewer buyers participate in the market when u increases. The buyer to seller ratio is 7.75 when $\delta = 0.88$ and $u = 1$, while the ratio is only 2.92 when $\delta = 0.88$ and $u = 8$. The total participation cost incurred by all buyers is the product of the average per-bidder cost and the ratio of buyers to sellers. As seen in Table 6, the participation costs drop by roughly 60% as u moves from 1 to 8 for the $\delta = 0.88$ case.

To place these results in the context of our finite model, one needs to recall that the limit model normalizes the measure of sellers to 1 and that there are (on average) 11.25 auctions per day in our data. To find the total cost incurred in the finite setting, we need to “de-normalize” the measure of sellers by multiplying the total participation cost figures by 11.25. This implies an ALPC in the finite market when $\delta = 0.88$ and $u = 1$ of \$782.78, which would drop to \$311.96 when $u = 8$.

5.2. Relative Importance of Platform Composition and Dynamic Incentives. Our model has two novel features relative to most of the empirical auctions literature: platform composition effects and dynamic incentive effects. Our goal in this section is to measure the relative importance of these two. As an illustrative example, we consider changes to the per-period participation cost κ . Aside from illuminating answers to questions of academic interest, this counterfactual provides practical guidance to eBay and other on-line market designers regarding what issues are of most importance when considering changes to a platform.

There are two effects when participation costs increase. First, agents’ continuation values drop, which in turn reduces demand shading and increases their bids. Holding the reserve price distribution G_R fixed, these *dynamic incentive* (DI) effects increase allocative efficiency since bids are now more likely to exceed the reserve price R . Second,

an increase of the participation cost drives low-value buyers out of the market, which reduces the buyer-seller ratio and strengthens the steady-state distribution of active bidder types. The consequences of buyer selection out of the market are referred to as *platform composition* (PC) effects.

We consider a range of participation costs from the estimated status-quo value, which we denote $\underline{\kappa} = \$0.0657$, through a maximum of \$10. Our goal is to separate the DI and PC effects, which are tied together intricately in equilibrium. For each counterfactual we consider the status-quo equilibrium with $\underline{\kappa}$ and replace either the bid and value functions (which drive the DI effect) or the buyer-seller ratio and bidder value distribution (which drive the PC effect) of an alternative equilibrium with $\kappa' > \underline{\kappa}$. The reader should keep in mind that neither of these exercises result in equilibrium outcomes; rather, they are meant to serve as a decomposition of the PC and DI effects.

To formally define the comparative statistics of interest, let $\mathcal{V}_\kappa(v)$ denote the value function for a bidder with value v in an equilibrium with participation cost κ . Let \mathcal{C}_κ denote the ratio of (active) buyers to sellers and λ_κ denote the matching parameter in an equilibrium with participation cost κ .⁴¹ Let f_{V_κ} and F_{V_κ} denote the analogous steady-state PDF and CDF of (active) bidder types and note that these live on support $[\underline{v}_\kappa, \bar{v}]$, with $\underline{v} < \underline{v}_{\kappa'}$ whenever $\underline{\kappa} < \kappa'$. Finally, let $\beta_\kappa(v)$ denote the equilibrium bidding strategy given participation cost κ . The probability of a buyer winning is:

$$\chi_\kappa(v; \mathcal{V}_\kappa, \lambda_\kappa, F_{V_\kappa}, \beta_\kappa) = G_R(\beta_\kappa(v)) \sum_{m=0}^{\infty} \pi_M(m, \lambda_\kappa) [F_{V_\kappa}(v)]^m.$$

The first term above captures the probability of an agent's bid exceeding the reserve price. The remaining terms are the probability that a buyer beats other competing bids. If all of the κ subscripts take on the same value, then χ_κ is generated by a steady-state SCE for that particular value of κ .

The allocative efficiency, \mathcal{W} , is a function of the endogenous variables considered:

$$\mathcal{W}(\mathcal{V}_\kappa, \lambda_\kappa, F_{V_\kappa}, \beta_\kappa) \equiv \mathcal{C}_\kappa \int_{\underline{v}}^{\bar{v}} s \chi_\kappa(s; \mathcal{V}_\kappa, \lambda_\kappa, F_{V_\kappa}, \beta_\kappa) f_{V_\kappa}(s) ds,$$

where for convenience we simply define $F_{V,\kappa}(v) = f_{V,\kappa}(v) = 0$ for each $v \in [\underline{v}, \underline{v}_\kappa]$.

Our metric for the role of DI effects in shaping welfare is the *dynamic gap*, defined by:

$$\mathcal{DG}(\underline{\kappa}, \kappa') \equiv \mathcal{W}(\mathcal{V}_{\kappa'}, \lambda_{\underline{\kappa}}, F_{V_{\underline{\kappa}}}, \beta_{\kappa'}) - \mathcal{W}(\mathcal{V}_{\underline{\kappa}}, \lambda_{\underline{\kappa}}, F_{V_{\underline{\kappa}}}, \beta_{\underline{\kappa}}).$$

⁴¹Including notation for both \mathcal{C}_κ and λ_κ is not necessary since the former can be computed from the latter.




FIGURE 8. Relative Size of Dynamic and Selection Gaps

The dynamic gap is computed by comparing equilibrium allocative efficiency generated by $\underline{\kappa}$ to an out-of-equilibrium market that uses the same matching parameter and steady-state distributions, but the value function ($\mathcal{V}_{\kappa'}$) and bidding strategy ($\beta_{\kappa'}$) from an SCE with a higher participation cost κ' . The idea is to hold fixed the endogenous quantities that correspond to PC effects (λ and F_V) while allowing DI effects (β and \mathcal{V}) to vary with κ .

The *platform gap*, \mathcal{PG} , captures the importance of PC effects in determining welfare:

$$\mathcal{PG}(\underline{\kappa}, \kappa') = \mathcal{W}(\mathcal{V}_{\kappa'}, \lambda_{\kappa'}, F_{V_{\kappa'}}, \beta_{\kappa'}) - \mathcal{W}(\mathcal{V}_{\underline{\kappa}}, \lambda_{\underline{\kappa}}, F_{V_{\underline{\kappa}}}, \beta_{\underline{\kappa}})$$

This gap is computed by comparing equilibrium allocative efficiency generated by $\underline{\kappa}$ to an out-of-equilibrium market with the same value function ($\mathcal{V}_{\underline{\kappa}}$) and bidding strategy ($\beta_{\underline{\kappa}}$) but matching parameters and steady-state distributions of an equilibrium with a higher cost κ' . Here we hold DI effects (β and \mathcal{V}) fixed and vary endogenous quantities that correspond to the PC effects (λ and F_V).

In Figure 8 we plot the ratio of the platform gap to the dynamic gap. When participation costs are low, the platform gap is only twice as large as the dynamic gap. However, as costs rise, the platform gap becomes as much as ten times larger than the dynamic gap. In short, it appears that understanding the platform composition effects of market changes is often many times more important than understanding the dynamic incentive effects of the changes.

5.3. Optimal Starting Prices. As has been regularly noted about the eBay marketplace, sellers tend to choose very low starting prices. In our data, almost 60% of the starting prices are set at the lowest possible value of \$0.99. It is easy to see that such a price is not optimal - a single seller could improve his profits if he set a starting price equal to \underline{v} ,

the lowest possible bid in the auction.⁴² From the perspective of a single seller, choosing the optimal starting price involves the same reasoning as in the classic optimal auctions literature: a high starting price can increase the final price paid by the winner, but it also risks that the good will go unsold.

We solve this problem numerically to get a sense of the strength of the incentives of the sellers to carefully choose a revenue maximizing starting price. For this exercise we assume that the seller has a supply cost of \$0. Since we are considering a deviation by a single seller in our limit game, the seller's deviation has no effect on market aggregates. As a result, we fix λ , F_V , and (e, β) at their status quo values. The problem the seller solves is:

$$\max_{r \geq 0} \Pr\{B^{(1)} \geq r\} E \left[\max\{r, B_M\} | B^{(1)} \geq r \right] \quad (46)$$

where $B^{(1)}$ is the highest bid in the auction and B_M is the highest competing bid.

Our results are remarkably stable across different choices of δ . The optimal starting price varies from a low of \$84.90 to a high of \$85.80. At the optimal starting price, the revenue generated is either \$122.30 or \$122.31 across all of the possible δ . This represents an increase in profit of just \$0.95 relative to a starting price of \$0.99.

The benefits from optimally choosing the starting price are small because each seller is matched with 7.96 bidders in expectation, which means that the competition between bidders is intense. Bulow and Klemperer [1996] show that in a static auction setting choosing the starting price optimally pales in comparison to adding a single extra bidder to the market.⁴³ With almost 8 bidders on average already participating, it should not be surprising that there is little room left for optimizing the starting price to have a significant effect on auction revenues.

6. APPROXIMATING A FINITE MODEL

We refer to a model with a finite number of buyers and sellers as a *finite model*. Since the real world is clearly finite, we ideally would have estimated and computed counterfactuals using a finite model. Moreover, one might be skeptical that our continuum model bears much relation to the finite model we would have liked to have worked with. The goal of this section is to both explain why we could not have conducted our analysis with a finite model and justify our use of the continuum model as an approximation of the more realistic finite setting. We first lay out the primitives of the finite model analog

⁴²Such a starting price would insure that if a single buyer was matched to the auction, the seller could extract some value from that buyer. It would have no effect if two or more buyers were matched to the auction as one of these buyers would necessarily set the sale price.

⁴³Since we do not provide a model of seller activity, we view Bulow and Klemperer [1996] as merely suggestive of what occurs in our setting.

to the continuum model described in Section 2, and we prove that a Bayesian Nash equilibrium of the finite model exists. We then prove that an SCE of the continuum model is an approximate equilibrium of the finite model. This approximation result is our justification for the use of a continuum model as a proxy for the intractable, but more realistic, finite model.

6.1. Primitives of the Finite Model. We consider a sequence of games indexed by N where N refers to the number of sellers that list goods for sale in each period. All variables pertaining to the N -agent game are superscripted with N .⁴⁴ Each seller has a reserve price R that is drawn randomly from the distribution G_R . The numbers of potential entrant buyers at $t = 0$ is denoted C_0^N . We assume that as $N \rightarrow \infty$

$$\frac{C_0^N}{N} \rightarrow C_0 \in \mathbb{R}_{++}.$$

The population of potential entrants in period t is C_t^N . Nature generates $\lceil N\mu \rceil$ new buyers at the end of each period and adds them to the set of potential entrants. Each time Nature generates a new potential entrant buyer, her private value v is drawn from T_V . The measure $F_{V,t}^N$ describes the distribution of potential entrant buyer values in period t of the N -agent game. As in the continuum game, buyers observe their own value for the good, the bidding cost, and the number and value distribution of the other potential entrant buyers in the game prior to choosing whether to enter. A bidder makes her choice of a bid without knowing either the number or identity of the other agents participating in the particular auction to which she is matched.

We now describe the stochastic matching process that assigns bidders to auctions in the finite setting. We denote the number of buyers that enter the market in period t as C_t^N . The buyers and sellers are randomly ordered into queues with the ordering independent across periods. Nature sequentially matches each seller in the respective queue with the next $k \in \{0, 1, \dots\}$ buyers from the buyer queue where k is a realization of random variable K that is distributed according to probability mass function $\pi(K; \lambda)$.

Intuitively (see formal proof in the technical appendix), if we consider a limit where the number of entering buyers and sellers grows without bound, then in the limit all of the entrants are matched into auctions. In the finite model, if the supply of entrants is not completely assigned to auctions, the unassigned buyers are referred to as *unmatched*

⁴⁴It is not difficult to allow for a random number of sellers in the finite game. If we denote the (potentially stochastic) number of sellers entering in period t of the N -agent game as S_t^N , we require that:

$$\frac{S_t^N}{N} \rightarrow 1 \text{ almost surely as } N \rightarrow \infty$$

We do not consider this extension due to the considerable number of notational aggravations it causes.

buyers. Unmatched buyers proceed to the next period without transacting. Conditional on being matched, a particular bidder wins her auction if her bid is larger than the maximum of all competing bids and the seller's reserve price. Ties between highest bidders are resolved by assigning the item to the tied bidders with equal probability, but if the highest bid is tied with the reserve price, then we assume the bidder wins the item.

Bidding strategies can be written as functions $\mathcal{O} : [0, 1] \times \mathbb{R}_+ \times \Delta([0, 1]) \times \Delta([0, 1]) \rightarrow [0, 1]$ with a typical bid denoted $\mathcal{O}(v_i, C_t^N, F_{V,t}^N, G_{R,t}^N)$. The entry decision for participating buyers is a function of the form $\theta : [0, 1] \times \mathbb{R}_+ \times \Delta([0, 1]) \times \Delta([0, 1]) \rightarrow \{Enter, Out\}$ with a typical realization $\theta(v_i, C_t^N, F_{V,t}^N, G_{R,t}^N)$. We let Σ denote the buyers' strategy space.

We use the notation $x^N(b, C_t^N, F_{V,t}^N, G_{R,t}^N) = 1$ (0) to denote the random event that a buyer wins (loses) an auction with a bid of b , and $p^N(b, C_t^N, F_{V,t}^N, G_{R,t}^N)$ denotes the random transfer from the buyer to the seller/eBay conditional on a bid of b . We also define:

$$\begin{aligned}\chi^N(b, C_t^N, F_{V,t}^N, G_{R,t}^N) &= E_t^N \left[x(b, C_t^N, F_{V,t}^N, G_{R,t}^N) \right] \\ \rho^N(b, C_t^N, F_{V,t}^N, G_{R,t}^N) &= E_t^N \left[p(b, C_t^N, F_{V,t}^N, G_{R,t}^N) \right]\end{aligned}$$

We superscript the expectation operator to emphasize that we are referring to the N -seller economy.

All agents discount future payoffs using a per-period discount factor $\delta \in (0, 1)$. The value function given a (symmetric) equilibrium strategy vector $\sigma = (\theta, \mathcal{O})$ for a bidder that chooses *Enter* is

$$\mathcal{V}^N(v, C_t^N, F_{V,t}^N, G_{R,t}^N | \sigma) = \chi^N v - \rho^N - \kappa + (1 - \chi^N) \delta E_t^N \left[\mathcal{V}^N(v, C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma) \right]$$

For a buyer that chooses *Out* we have

$$\mathcal{V}^N(v, C_t^N, F_{V,t}^N, G_{R,t}^N | \sigma) = \delta E_t^N \left[\mathcal{V}^N(v, C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma) \right]$$

We use the notation $\mathcal{V}^N(v, C_t^N, F_{V,t}^N, G_{R,t}^N | \sigma'_i, \sigma_{-i})$ when buyer i uses strategy σ'_i and all other agents follow σ .

We use the following definition of an equilibrium in our finite games.

Definition 6.1. The strategy vector $\sigma = (\theta, \mathcal{O})$ and the initial state $C_0^N \in \Omega_N$ and $F_{V,0}^N, G_{R,0}^N \in \Delta_N([0, 1])$ is an ε -Bayes-Nash Equilibrium (ε -BNE) of the N -agent game if for all bidder values v we have

$$\text{For all } \sigma'_i \in \Sigma_C, \mathcal{V}^N(v_i, C_0^N, F_{V,0}^N, G_{R,0}^N | \sigma) + \varepsilon \geq \mathcal{V}^N(v_i, C_0^N, F_{V,0}^N, G_{R,0}^N | \sigma'_i, \sigma_{-i})$$

Given the dynamic nature of our game, a solution concept that incorporates some notion of perfection might be expected. Consider the two ways in which an ε -BNE can

yield an $\varepsilon > 0$. First, it may be that the agent does not exactly optimize with respect to high probability events, which results in a small loss with high probability. Second, the strategy may not optimize with respect to very rare events. Failing to optimize with respect to rare events can be approximately optimal but severely violate perfection. A stationary strategy can be an ε -BNE even though perfection may not even be approximately satisfied at the histories of the finite game in which the market aggregates differ significantly from the stationary state.

We now prove that there exists an exact BNE for our finite model if there exists an equilibrium for the static spot market (i.e., when $\delta = 0$). Our proof is constructive in the sense that it uses the equilibrium of the static version of the model to solve for an equilibrium of our dynamic model. For those interested in applying our work in other settings, this is useful since equilibrium existence in static auctions has been established for a wide array of pricing rules. From a theoretical perspective, it is interesting to note that an equilibrium of a static model can be easily mapped into an equilibrium of our dynamic model. The key insight is that each agent's effective valuation from the perspective of today is her private value minus the opportunity cost of winning.

Proposition 6.2. *Suppose that if $\delta = 0$ there exists an equilibrium $\tilde{\sigma} = (\tilde{\theta}, \tilde{\mathcal{O}})$. Then we can define the equilibrium $\sigma = (\theta, \mathcal{O})$ when $\delta > 0$ as*

$$\begin{aligned}\theta(v, C_t^N, F_{V,t}^N, G_{R,t}^N) &= \tilde{\theta}(v - \delta E_t^N [\mathcal{V}^N(v, C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma)], C_t^N, F_{V,t}^N, G_{R,t}^N) \\ \mathcal{O}(v, C_t^N, F_{V,t}^N, G_{R,t}^N) &= \tilde{\mathcal{O}}(v - \delta E_t^N [\mathcal{V}^N(v, C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma)], C_t^N, F_{V,t}^N, G_{R,t}^N)\end{aligned}$$

6.2. Approximating the Large Finite Model. It is not difficult to see why the model becomes too computationally complex to solve precisely as $N \rightarrow \infty$. In the N -agent game, the bidder's strategy must condition on all possible values of C_t^N , $F_{V,t}^N$, and $G_{R,t}^N$, which means the complexity of the strategies grows exponentially with N . The bidding strategy in the continuum need only condition on the values of C_t , $F_{V,t}$, and $G_{R,t}$, which evolve deterministically in equilibrium.

Moreover, computing the stochastic evolution of the state variables in the N -agent game is difficult. For example, suppose a bidder knows the type distribution for the current period, $F_{V,t}^N$, an N -dimensional step function. In order to precisely compute the type distribution for the following period, she must first gather a large amount of information on agents exiting the market. After conditioning on the set of buyers being removed from the game, she must also account for the infusion of $\lceil N\mu \rceil$ new potential buyers with private valuations being drawn from T_V . This requires computing probabilities that tomorrow's CDF of types, usually an object with more than N dimensions

as there are often many more buyers than sellers, and then condensing this information into one single distribution, $F_{V,t+1}$.

Our goal is to prove that the limit model approximates the large finite model. The foundation of our proof is a mean field result that proves that the evolution of the continuum game and the economy of a finite game with sufficiently many players are approximately the same over finite horizons. Mean field results usually require strong continuity conditions on the evolution of the economic primitives and on the strategies adopted by the agents, conditions that we need to prove hold despite auction models admitting a wide array of possible discontinuities. In addition, since the within-period matching process of the finite game samples without replacement from a finite set of buyers, there is a nonzero correlation between bidder values across auctions that close within the same period. We prove that as the market grows, the auctions become independent of one another. In addition, there is also a positive probability that a positive mass of buyers is unmatched, and we show that the fraction of unmatched buyers vanishes as the size of the market increases.

With our mean field result in hand, we demonstrate that the expected buyer utility in the large finite game and the limit game are approximately the same. This insight translates into our approximation result, which proves that any exact SCE strategy of the limit game is an ε -BNE of the finite game with sufficiently many players.

Proposition 6.3. *Consider a SCE (σ, C, F_V, G_R) and assume $e(C, F_V, G_R) < 1$. For any $\varepsilon > 0$ we can choose N^* and $\eta > 0$ such that for all $N > N^*$, σ is an ε -BNE strategy if $(\omega_0^N, F_{V,0}^N, G_{R,0}^N)$ satisfies*

$$\|C_0^N - C\| + \|F_{V,0}^N - F_V\| + \|G_{R,0}^N - G_R\| < \eta \quad (47)$$

Proposition 6.3 may be seen as providing an approximation to the actual equilibrium being played within the data-generating process, but it admits an alternative interpretation as a behavioral strategy. If one assumes that agents are subject to small computational costs, then in large markets it may be that they follow SCE behavioral predictions in lieu of solving a complex optimization problem for a vanishing benefit. Finally, note that while our result requires that the aggregate states be close in period 0, if we assume that seller and bidder types are drawn from F_V and G_R with numbers close to N_C and N_B , then $(C_0^N, F_{V,0}^N, G_{R,0}^N) \rightarrow (C, F_V, G_R)$ almost surely as $N \rightarrow \infty$. In other words, Equation 47 above is very likely to hold in large markets, and becomes increasingly so as $N \rightarrow \infty$.

7. CONCLUSION

Our goal has been to provide a model of a dynamic auction platform that is both rich enough to capture the salient features of the market (e.g., the large number of auctions concluding each day, the cost of participation) and yet remain tractable enough to facilitate empirical analysis. To accomplish this, we have developed a model with a continuum of buyers and sellers that is easy to estimate and solve, and we have shown that this model approximates the more realistic setting with a finite number of agents. We have also demonstrated that the structural components of this model can be identified from observables that are commonly available from platform markets. In constructing these identification results we have overcome several important problems including sample selection in the number of spot-market competitors and allowing for pricing rules that give rise to static demand shading incentives. Finally, we have also proposed a simple but flexible GMM estimator for the structural primitives.

Most platform markets exist in order to eliminate barriers to trade and allow for buyers and sellers to interact in a relatively low-friction environment. However, the sheer size of the markets may give rise to search frictions which prevent market outcomes from attaining the social optimum. We have estimated our model within the context of the market for Kindle Fire tablets, and we use these estimates both to compute the welfare loss under the present design and to suggest novel designs to mitigate these welfare losses. We begin by providing a “model-anemic” analysis that relies on the observed bid distribution. We find that at least 28.47% of the auctions close with a highest losing bidder that ought to be allocated the good in an efficient within-period allocation.

We then use our structural estimates to put a value on the deadweight loss and to study alternative spot market mechanisms that might eliminate (some of) the welfare loss due to search frictions. We find that over 36% of the auctions end with an inefficient allocation, and a 13.5% welfare loss can be attributed to the decentralized nature of the mechanism. This outcome implies that the single-unit auction market attains three quarters of total possible welfare improvement over a pure lottery system. By taking small steps toward a more centralized market structure— such as running multi-unit, uniform-price auctions with as few as 4 units each—2/3 of the welfare loss can be recovered.

Another conclusion of our analysis relates to the importance of intertemporal incentives. In online auction markets, bid shading is driven by the opportunity cost of winning today, which depends on three main factors: market tightness (ratio of buyers to sellers), market composition (ratio of high-value buyers to low-value buyers), and time preferences. The dynamic incentive to shade one’s bid is much larger in magnitude than the more commonly studied static bid shading incentives generated by nontruthful pricing

mechanisms. In other words, it is more important for buyers to understand their opportunity costs than it is for them to understand how to strategically respond to nontruthful pricing mechanisms.

We attempt to disentangle the welfare effects of dynamic incentives, which is the primary source of bid shading, from the platform composition effects governing the selection of buyer types into the market, which governs the steady-state distribution of buyer values and the buyer-seller ratio. We consider different participation costs, and we compute the magnitude of the welfare effects (relative to the status quo) of the dynamic incentive and the platform effects. We find that the platform composition effects are at least twice as important as the dynamic incentive effects. Our primary takeaway is that understanding endogenous selection into the market is critically important for judging the effects of possible mechanism changes.

In future work we hope to estimate a structural model of the sellers' actions on the eBay market platform. In a contemporaneous paper, we are estimating the value of sellers of the Kindle product within the posted price Buy It Now market on eBay. The posted price framework gives sellers a strong incentive to carefully balance the trade-off between price and probability of sale, which makes the resulting estimates of seller reservation values plausible. By integrating the estimates of bidder values from the auctions with seller reservation values from the posted price setting, we hope to be able to derive the optimal participation fee schedule for a profit-maximizing platform designer like eBay, and the related welfare implications from the social planner's perspective. However, until a credible estimate of seller values is available, these interesting and important questions remain elusive.

REFERENCES

- Susan Athey and Philip A. Haile. Identification of standard auction models. *Econometrica*, 70 (6):2107–2140, November 2002.
- Ned Augenblick, Muriel Niederle, and Charles Sprenger. Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics*, forthcoming, 2016.
- Eduardo M. Azevedo and Jacob D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, forthcoming, 2016.
- Matthew Backus and Gregory Lewis. A demand system for a dynamic auction market with directed search. *typescript, Cornell University*, 2012.
- Matthew Backus and Gregory Lewis. Dynamic demand estimation in auction markets. *NBER Working Paper*, No. 22375, 2016.

- Patrick Bajari and Ali Hortaçsu. The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34:329–55, 2003.
- Eric Budish. Sequencing and information revelation in auctions for imperfect substitutes: Understanding ebay's market design. *typescript, Harvard University*, 2008.
- Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *American Economic Review*, 86 (1):180–194, March 1996.
- Stephen Burks, Jeffrey Carpenter, Lorenz Götte, and Aldo Rustichini. Which measures of time preference best predict outcomes: Evidence from a large-scale field experiment. *Journal of Economic Behavior & Organization*, 84:308–320, 2012.
- Xiaohong Chen. Large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics*, 6(76):5549–5632, 2007.
- Dominic Coey, Bradley Larsen, and Brennan C. Platt. A theory of bidding dynamics and deadlines in online retail. *Typescript, Stanford University Economics Department*, 2016.
- P. C. Consul and G. C. Jain. A generalization of the poisson distribution. *Technometrics*, 15(4):791–799, 1973.
- Carl de Boor. *A Practical Guide to Splines, Revised Edition*. Springer-Verlag: New York, 2001.
- Liran Einav, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7, 2015.
- Drew Fudenberg, David Levine, and Wolfgang Pesendorfer. When are nonanonymous players negligible? *Journal of Economic Theory*, 79(1):46–71, 1998.
- Edward J. Green. Noncooperative price taking in large dynamic markets. *Journal of Economic Theory*, 22(2):155–182, 1980.
- Edward J. Green. Continuum and finite-player noncooperative models of competition. *Econometrica*, 52:975–993, 1984.
- Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68:525–574, 2000.
- Fumio Hayashi. *Econometrics*. Princeton University Press, 2000.
- Brent R. Hickman. On the pricing rule in electronic auctions. *International Journal of Industrial Organization*, 5:423–433, 2010.
- Brent R. Hickman and Timothy P. Hubbard. Replacing sample trimming with boundary correction in nonparametric estimation of first-price auctions. *Journal of Applied Econometrics*, 30(5):739–762, 2015.
- Brent R. Hickman, Timothy P. Hubbard, and Harry J. Paarsch. Identification and estimation of a bidding model for electronic auctions. *Quantitative Economics, forthcoming*,

2016.

- Werner Hildenbrand. *Core and Equilibria of a Large Economy*. New Jersey: Princeton University Press, 1974.
- Hugo Hopenhayn and Maryam Saeedi. Dynamic bidding in second-price auctions. *type-script, University of California–Los Angeles Economics Department*, 2016.
- Matthew O. Jackson and Alejandro M. Manelli. Approximately competitive equilibria in large finite economies. *Journal of Economic Theory*, 77(2):354–376, 1997.
- R. J. Karunamuni and S. Zhang. Some improvements on a boundary corrected kernel density estimator. *Statistics and Probability Letters*, 78:499–507, 2008.
- Fuhito Kojima and Parag A. Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99(3):608–627, 2009.
- Elena Krasnokutskaya. Identification and estimation in procurement auctions under unobserved auction heterogeneity. *Review of Economic Studies*, 28, 2011.
- Iyer Krishnamurthy, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *Management Science*, 60(12):2949–2970, 2014.
- Richard MacLean and Andrew Postlewaite. Informational size and incentive compatibility. *Econometrica*, 70(6):2421–2453, 2002.
- D.R. McDonald. The local limit theorem: A historical perspective. *Journal of the Iranian Statistical Society*, 4:73–86, 2005.
- Roger B. Myerson. Population uncertainty and poisson games. *International Journal of Game Theory*, 27:375–392, 1998.
- Denis Nekipelov. Entry deterrence and learning prevention on ebay. *mimeo*, 2007.
- Yoshihiko Otani and Joseph Sicilian. Limit properties of equilibrium allocations of walrasian strategic games. *Journal of Economic Theory*, 51(2):295–312, 1990.
- Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 2:601 – 653, 2014.
- Michael Peters and Sergei Severinov. Internet auctions with many traders. *Journal of Economic Theory*, 130:220–245, 2006.
- Brennan C. Platt. Inferring ascending auction participation from observed bidders. *Type-script, Brigham Young University Economics Department*, 2015.
- James Roberts. Unobserved heterogeneity and reserve prices in auctions. *RAND Journal of Economics*, 44:712–732, 2013.
- John Roberts and Andrew Postlewaite. The incentives for price-taking behavior in large exchange economies. *Econometrica*, 44(1):112–127, 1976.
- Alvin E. Roth and Axel Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *American Economic Review*, 92 (4):1093–1103, 2002.

Hamid Sabourian. Anonymous repeated games with a large number of players and random outcomes. *Journal of Economic Theory*, 51(1):92–110, 1990.

Gabriel Weintraub, Lanier Benkard, and Benjamin Van Roy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.