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 large number 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. Solving our model with a finite number of bidders is computationally intractable due to the curse of dimensionality, so we prove that a continuum version of our model provides a good approximation of an equilibrium in the finite model. We use the approximation to 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 13.5% of total possible market surplus. We also find that partial centralization - for example, runnng half as many 2-unit, uniform price auctions each day - would eliminate a large fraction of the inefficiency, but yield lower seller revenues. Our results highlight the importance of understanding platform composition effects—selection of agents into the market—in assessing the implications of market design.

Keywords: Dynamic Auctions, Approximate Equilibrium, Internet Markets.

JEL Codes: C73, D4, L1

Date: August, 2016.

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. Boehnke gratefully acknowledges the support of the Center of Mathematical Sciences and Applications at Harvard University during the course of this work.

*Email: acreed@berkeley.edu
§Email: boehnke@cmsa.fas.harvard.edu
†Email: hickmanbr@uchicago.edu.

1. INTRODUCTION

Market platforms are increasingly important in today's economy. Familiar examples include platforms for retail goods like eBay, StubHub, Christie's, Sotheby's, and Manheim, but more recently, internet technology has even given rise to platform markets for freelance labor, like Upwork. These firms are all major players in their respective economic spheres. American car dealers sourced 30% of their used inventory from auction houses in 2007 (see Roberts [2013]). In 2014 eBay reported USD$82.95 billion in sales volume and 8.5% annual growth after nearly two decades in business.[1] Christie's, Sotheby's, StubHub, Manheim, and Upwork each host annual transaction volumes in the billions or hundreds of millions of dollars. These firms each have a distinguishing characteristic in common: none of them engage as active participants in the transactions they host, none directly buy or sell services, and none maintain any inventories of their own goods. In other words, each firm's platform exists solely for the purpose of creating a forum where buyers and sellers can easily match for trade.

While much research has focused on static auctions, much less is known about the role of market design in shaping bidding incentives, welfare, and revenues in dynamic platform markets. On a platform, a large number of buyers and sellers participate in essentially simultaneous auctions each period, and agents know that if they are unsuccessful in consummating a trade today, they can return to the market in future periods to try again. In this paper we provide a rich model of such an auction platform in which buyers are matched to the sellers' auction listings each period, auction winners (and the associated sellers) exit the market, and new players enter each period. We include a costly endogenous entry decision, and endogenous bidding choice to capture the time and effort costs of participation.

Our rich model captures the salient features of platform markets and allows for a broad class of spot-market pricing mechanisms within individual listings. The model characterizes market evolution over time, but a curse of dimensionality arises from an intractably large state space when even moderately large numbers of agents participate. This problem creates an important barrier to empirical work, and to solve it we also develop a more tractable version of the model with a continuum of buyers and sellers. Our first contribution is proving that the equilibria of the continuum model provide a reasonable approximation of equilibria of the finite model, and the continuum model can therefore be used as a basis for empirical work. Our formal justification of the continuum model as an approximation to the finite one is, to the best our knowledge, a novel contribution within the platform markets literature.

[1]Information downloaded from
https:investors.ebayinc.comsecfiling.cfm?filingID=1065088-15-54&CIK=1065088 on 11/17/2015.

FIGURE 1. Static Versus Dynamic Demand Shading Incentives

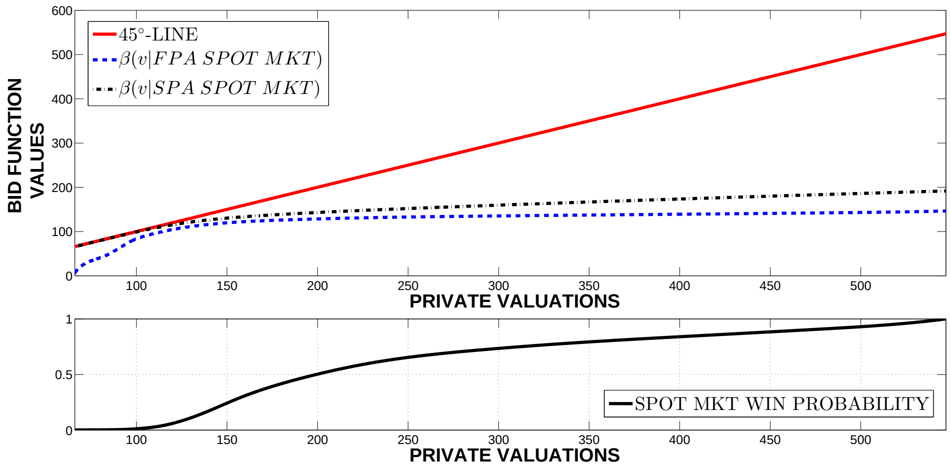

We use our limit approximation to build an empirical model of dynamic bidding behavior within a homogeneous goods market hosted on eBay. Market dynamics produce incentives for demand shading (*i.e.*, bidding strictly below one's private valuation) even when the spot-market pricing rule has the familiar second-price form. Intuitively, demand shading arises from the idea that if the hypothetical spot-market price rose too high, the bidder would prefer to forgo a purchasing opportunity today (even one below her private valuation) in expectation of paying a lower price tomorrow. This opportunity cost determines the degree of demand shading: if market conditions tomorrow are expected to be favorable to buyers, then incentives to bid aggressively today are weakened. Given a value for the time discount factor, this demand shading factor is nonparametrically identified from observables that are readily available on many platform markets like eBay. We also show that when the spot-market pricing rule is non-second-price, the strategic incentive to shade bids present in static settings is layered on top of the dynamic incentives in a parsimonious way. This model extension 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. We outline conditions under which the static and dynamic demand shading factors can be separately identified.

Interestingly, our empirical analysis reveals that the degree of bid shading incentivized through intertemporal opportunity costs is significantly larger than demand shading generated by the use of a non-truthful spot-market auction mechanism. Figure 1, generated using our structural estimates, demonstrates this point. In the upper pane the 45°-line (solid) is plotted against the dynamic bidding strategy resulting from the two most commonly used spot-market mechanisms: second price (dash-dot line) and first price (dashed line). The lower pane plots the win probability for context. The difference between the dash-dot line and the dashed line represents the influence of static incentives for shaping bidding 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 dynamic demand shading factor is an order of magnitude larger than the static demand shading factor. To the best of our knowledge, ours is the first study to highlight the importance of dynamic demand shading incentives relative to the traditional demand shading incentives present in static auctions.

Decentralized auction platforms like eBay facilitate the interaction of a set of heterogeneous buyers and sellers on opposite sides of the market. Agents on each side of the market have preferences over the number of agents present on the other side of the market. For example, buyers would like the platform to be used by many sellers and few buyers to reduce bidding competition in the auctions. In our setting the agents also have preferences over the types of agents on the other side of the market. For example, sellers would prefer the buyers to have high values since this would increase the bids in the auctions. The fact that agents care both about the number and type of agents on both sides of the market makes our model more complex than many commonly used in the platform economics literature. In addition, our model raises new and interesting issues when it comes to, for example, analyzing the effects of participation costs on the endogenous composition of platform market participants.

Finally, the dynamic demand shading incentives and platform composition effects interact in complex ways. For example, any change in the market structure that increases the continuation values of the average buyer (*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 their probability of winning an auction drops. These countervailing effects make it unclear whether a platform must necessarily benefit from improvements in market efficiency.

We use our model and an extensive dataset on new Amazon Kindle Fire tablets to estimate structural model primitives including the buyer-seller matching process, the monetized cost of participation, and the steady-state distributions of buyer valuations and seller reserve prices. The inclusion of a per-period participation cost—representing the value of foregone time and effort to find a listing to bid on—turns out to be an important regulator of the number and types of buyers in the market. While the participation cost we find is low, on the order of $0.10, it has a non-trivial impact on low-value buyers who must repeatedly bid in order to win an item.

A key feature of our proposed estimator is that it requires only observables that are readily available on many platform websites. In particular, if bid submission times are randomly ordered, then some auction participants watching an item with 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 accounts for this source of sample selection, and requires only that we observe this lower bound on auction participation, the seller reserve price, and the highest losing bid.

Having estimated the economic primitives of our model, we perform three comparative statics exercises. Our first one investigates the efficiency of the decentralized auction platform. Even when the spot-market mechanism is known to allocate efficiently *within a given auction*—where we adopt the usual notion of efficiency as the tendency for goods to be allocated to those who value them most—platforms like eBay still exhibit search frictions due to their sheer size, which leads to inefficiency. To fix ideas, consider a very simple platform market where each period there are two identical auctions and four bidders participating; two with high private values and two with low values. The social planner concerned with allocative efficiency would prefer that one high type and one low type be matched to each auction, thereby ensuring that (in any monotone bidding equilibrium) both units of the good will go to high-value buyers. However, because of private information and random matching, there is always positive probability that one of two things will go wrong within a given auction. First, a high-value buyer that ought to receive the good in an efficient allocation may end up competing against another high-value buyer, which means that one of them will not receive a unit of 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 an efficient outcome.

We measure the inefficiency in various ways. As a robustness check, we conduct what we refer to as a "model anemic" analysis that employs only our raw data and estimates of the buyer-seller matching process to place a lower bound on the frequency of inefficiency with as few assumptions as possible. We find that at least 28% of auctions result in inefficient allocations. Next, in order to identify the precise frequency and deadweight loss of inefficient allocations on eBay, we take additional measurements using our full structural model. We find that 36% of auction listings allocate inefficiently, leading to a deadweight loss of roughly 14% of total possible gains from trade. We then compute a series of counterfactuals to study alternative platform designs that are increasingly "centralized," meaning that eBay hosts the same number of items for sale using fewer multiple unit auctions. A fully centralized mechanism would be a single multi-unit auction that clears the market for all goods posted by sellers once per day. We find that most of the welfare gains from centralization can be recovered by moving from single-unit auctions to 2- or 3-unit auctions.

As for sellers, one might naturally assume that increased efficiency would lead to increased revenue as well. However, this turns out not to be the case because of two effects. First, *platform composition* (PC) effects alter the profile of buyers in the market as low-value buyers exit due to a drop in their probability of winning. Second, *dynamic incentive* (DI) effects (*i.e.*, the demand shading driven by opportunity costs) alters the bidding behavior of the remaining agents. Empirically, the net effect is that centralization results in a small *loss* of revenue to sellers.

Our second counterfactual exercise considers the relative importance of PC and DI effects for market efficiency. To this end, we consider an increase of the participation costs for bidders from our estimated level of \$0.0657 up to a value of \$10.00. The changing cost generates PC effects on the steady-state ratio of buyers to sellers and the distribution of buyer types in the market by altering the cutoff type of buyer that is just indifferent to participating on the platform. The participation cost change also produces DI effects by influencing opportunity costs since entering the market next period costs more when κ increases. We compare the relative magnitudes of these two effects through a counterfactual decomposition, and we find that PC effects are between two and ten times more important than the DI effects for driving market efficiency.

The final comparative static we consider is an analysis of seller behavior. Over half of sellers in our data set reserve prices at \$0.99, the default reserve price on eBay.[2] Within a dynamic platform market, optimal reserve price calculations diverge from classic mechanism design results by Myerson [1981] since sellers may also return to the market in a future period to re-list their good if it does not sell today. We show that, under our point estimates, the most profitable reserve price given a supply cost of \$0 is roughly \$85, but the benefit relative to a reserve price of \$0 is less than \$1. The extremely small benefit we estimate resolves the puzzle of why sellers typically choose not behave strategically. Our second and third counterfactual exercises together provide the first dynamic platform characterization analogous to a classic result by Bulow and Klemperer [1996] within static auctions: although it is possible to use market design to shape bidding behavior, participation and market composition are more important concerns.

The remainder of this paper has the following structure. In Sections 2 and 3 we develop a theory of bidding within a dynamic platform, and we establish our large market

[2]On eBay.com, the term for what we call *reserve price* here is actually "starting price," and the term "reserve price," as employed by eBay, refers to a hidden starting price. However, since the term *reserve price* is used more commonly among economists for what eBay calls a starting price, we maintain the traditional academic parlance to avoid confusion among our target audience.

approximation result. In Section 4 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 5 we present our model estimates. Finally, section 6 presents our counterfactuals on welfare, the relative importance of selection and dynamic incentives, and optimal reserve prices. Most of the proofs are relegated to an online technical appendix.

1.1. **Related Literature.** Ours is not the first paper to study a dynamic bidding model where intertemporal opportunity costs shape behavior. The earliest structural work in this vein, Jofre-Bonet and Pesendorfer [2003], focuses on firms' time-varying production capacity constraints (work backlogs) in repeated procurement auctions, with more recent extensions by Balat [2013]. Platform markets, by contrast, typically host sales for consumer products where bidders have single-unit demands and are thought to permanently exit the market after transacting once. Several other papers study such markets, including Zeithammer [2006], Said [2011], Coey, Larsen, and Platt [2016], and Backus and Lewis [2016], with the latter two papers being most closely related to ours.

Coey et al. develop a model of eBay with the goal of understanding why auctions and fixed-price listings both co-exist on the same platform. Their model is agnostic about buyer consumption utility, and all demand-side heterogeneity is driven by hard deadlines by which buyers must transact. The model of Backus and Lewis is most similar to ours. Their focus is on estimating a rich demand model with heterogeneous buyer preferences and flexible substitution patterns for distinct goods, while our model depicts heterogeneous buyer preferences for homogeneous goods. Given the complexity of the type space they study, their identification strategy naturally requires a richer set of observables than ours. In contrast, our paper takes a different focus from the above two papers by studying market composition (selection of buyers into the market), the efficiency of market allocations, and evaluating a set of market redesign counterfactuals. Because of the various modeling choices driven by distinct sets of research questions, we view these three papers as complementary to one another.

Our paper makes several new contributions, both methodological and applied. First, empirical work necessarily focuses on markets with finitely many agents, but existing models (including ours) employ a continuum steady-state model for tractability. We provide rigorous micro-foundations to this simplification by starting with a more realistic (though intractable) finite model, and then we prove that the continuum steady-state model is a limiting approximation. Second, our foundational theory and identification argument are both fully non-parametric aside from basic regularity conditions. Much existing empirical work on eBay either assumes that all bids within each auction are directly observed by the econometrician or resorts to parametric assumptions (*e.g.*, Poisson distributed auction sizes) in order to correct for problems of limited observables (*e.g.*, when some bidder identities are unobserved). Our identification argument corrects for limited observables without imposing parametric forms. Our empirical implementation provides strong evidence for the benefits of added flexibility in the model.

Third, our model allows for multiple sources of demand shading incentives by accommodating a variety of spot-market pricing mechanisms. Aside from providing generality, our model neatly decomposes dynamic demand shading incentives (intertemporal opportunity costs) from static demand shading incentives (strategic calculations specific to spot-market competition). As it turns out the former are much more important than the latter (see Figure 1); to the best of our knowledge we are the first to uncover this empirical insight. Fourth, our model contains an explicit search cost among the set of structural primitives. This allows us to monetize wasted resources due to time and effort involved in search and helps explain buyer selection into the market. Fifth, we explore a novel set of market re-design experiments that allow us to characterize counterfactual steady states under alternative, more efficient market structures. These counterfactuals illuminate the various mechanisms driving allocative inefficiency on eBay while controlling for changes in market buyer composition and bidding behavior.

Our theoretical large-market approximation results belong to a larger body of research which we briefly survey here. Early papers focused on conditions under which underlying game-theoretic models could be used as strategic micro-foundations 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]).

Another related paper is Hickman [2010] that shows that the pricing rule on eBay is actually a hybrid of a first-price and a second-price mechanism. Hickman, Hubbard, and Paarsch [2016] explore the empirical implications of the non-standard pricing rule on eBay within a static, one-shot auction model with no binding reserve prices. We build on these two papers in the following ways. First, our model incorporates both dynamic and static demand shading incentives. Second, we extend 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

Here we develop our model of a dynamic platform market; we will state important results here, but the associated proofs, being fairly lengthy, are left to a technical appendix. Both buyers and sellers participate in platform markets, but our focus will be on the strategic choices of the buyers within the market. We treat sellers collectively as a source of exogenous supply, and we take their decisions (*e.g.*, entry/exit and reserve prices) as exogenous and fixed within the model.[3] This modeling choice is driven by the fact that observables tied to the seller side of our the eBay market are less amenable to econometric identification, making empirical work a challenge.

The market evolves in discrete time with periods indexed $t \in \{0, 1, 2, ...\}$. We refer to the set of buyers present at the start of period t as *potential entrants*; each period they make decisions based on the observed number and type distribution of the other potential entrants and their own types. The first choice each potential entrant must make in a period 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.

After choosing *Enter*, each entrant formulates a strategic bid, knowing that she will be randomly matched to an auction listing where the number and types of other entrants matched to that particular auction will be unobservable to her. The form of the random matching process, the steady-state 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.[4] The auctions are resolved using the relevant price setting mechanism (*e.g.*, a first-price or second-price rule). If a buyer wins the spot-market

[3] As we discuss in greater length in section 6.3, given a production cost of \$0, the sellers earn less than \$1 in increased revenue by moving from a reserve price of \$0 to the revenue maximizing reserve price.

[4] The 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 et al. [2016], and Backus and Lewis [2016]. We find evidence in our empirical exercise (see Figure 1) and Section 5 below) consistent with the idea that incentives tied to inter-auction dynamics occupy a predominant role in bidders' payoffs.

auction, then she exits the game at the end of the period. All surviving buyers remain in the market for the next period. We summarize the timing in Figure 2.

FIGURE 2. Timing within a period

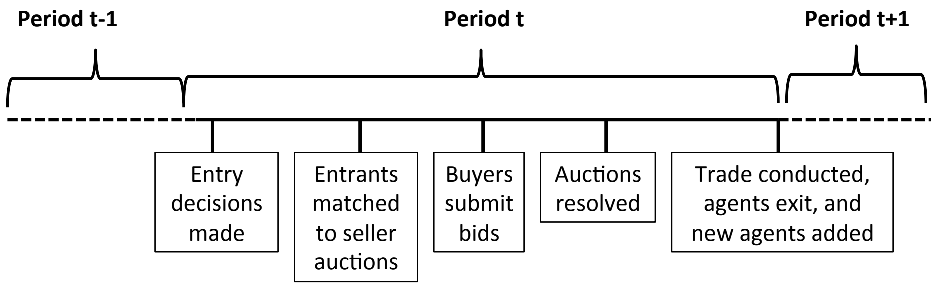

With an eye toward facilitating empirical work, several challenges must be addressed. Real-world platforms involve large but finite numbers of buyers and sellers, but computing equilibria for such markets is intractable. Therefore, our strategy here is to develop a simplified version of reality with a model having infinitely many economic agents. We then prove that this simplification is empirically relevant in the following sense: a bidder's dynamic value function under the limit 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 suggests that the tractable continuum model can be used in empirical work on platform markets with a large number of participants. Although our primary focus here is empirical, it is also worth noting that this result provides a novel theoretical contribution to the large markets literature as discussed in the previous section. 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 the analysis of platform markets.[5]

2.1. **Model Primitives.** We index buyers with i, and the value that buyer i places on the product is denoted v_i. We assume throughout that buyers have demand for a single unit, and the buyers' values for the good do not change with 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

[5]Part of our definition of a "well-behaved" pricing rule is one that only allocates the object to the highest bidder with positive probability, which we implicitly assume moving forward. In our empirical application we consider data from eBay, which uses an unusual hybrid pricing rule combining elements of both second-price and first-price auctions. See Section 2.1 for further discussion.

fee that the platform designer imposes. 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. Moving forward, we will use the terms "bidding cost" and "participation cost" interchangeably in reference to the parameter κ.

In our limit results we consider a sequence of games indexed by N. All of the games have a finite number of buyers and sellers, and we will often refer to the N-agent game. The variable N scales the number of buyers in the game at $t = 0$ and the number of new buyers entering the game each period. The reader should note that N is an index of size and does not imply that game has precisely N players.

All variables pertaining to the N-agent game are superscripted with N. Each period $S_t^N = \lceil NS^\infty \rceil$ sellers have goods for sale, where S^∞ is the measure of sellers each period of the continuum game. Each seller has a reserve price R that is drawn randomly from the distribution G_R.[6] G_R may have a mass point, but only at the lowest possible reserve price, $r = 0$, and has a density $g_R(R|R > 0)$ that is strictly bounded away from zero over the rest of its support $(0, \bar{r}]$. The realized distribution of reserve prices in the N-agent game in period t is denoted $G_{R,t}^N$. The numbers of potential entrant buyers at $t = 0$ is denoted C_0^N.[7] We assume that as $N \to \infty$

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

The population of potential entrants in period t is C_t^N. The parameter μ scales the number of new potential entrants added at the end of each period relative to the number of potential entrants present in period 0. In the N-agent game, Nature generates $\lceil N\mu \rceil$ new potential entrants at the end of each period and adds them to the set of potential entrants. We use the variable $\boldsymbol{\omega}_t^N = \left(\frac{C_t^N}{N}, \frac{S_t^N}{N} \right)$ to track the number of buyers and sellers at the beginning of period t. $\Omega_N \subset \Omega$ denotes the set of $\boldsymbol{\omega} \in \Omega$ compatible with the definition of the N-agent game.

Each time Nature generates a new potential entrant buyer, her private value v is drawn from an atomless cumulative density function (CDF) $T_V(\cdot)$ with probability density function (PDF) $t_V(\cdot)$. We assume that t_V is strictly positive over the support $[0, 1]$. The measure $F_{V,t}^N$ describes the distribution of potential entrant values at the beginning of period t of the N-agent game, including both newly generated potential entrants and ones remaining from period $(t-1)$, and $F_{V,t}^N$ is an element of the space of probability

[6] We use the letter G to refer to the cumulative density functions (CDF) 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.

[7] Since the letter B is used later on to denote bids, we chose C for "consumer" to represent the number of buyers flowing into the market each period.

measures over $[0,1]$, denoted $\Delta([0,1])$. Unless stated otherwise, $\Delta([0,1])$ is endowed with the weak-* topology. We let $\Delta_N([0,1])$ denote the set of empirical probability measures that can be generated by N realizations of $[0,1]$ and endow this space with the relative topology inherited from $\Delta([0,1])$.

Before deciding whether to enter, a potential entrant observes her own value for the good, the bidding cost, and the number and value distribution of the other potential entrants in the game. An entrant makes her choice of a bid without knowing either the number or identity of the other agents participating in the particular auction to which he or she is matched. We now describe the stochastic matching process that assigns bidders to auctions. Our formalization is designed to capture the important aspects of a scenario where bidders randomly arrive at the platform market (*e.g.*, the eBay homepage) at some point during the day, search for the desired product, and bid on one of the items that they happen to see (*e.g.*, an auction listing near the top of the first page of eBay's keyword search results). We implicitly assume that the buyers bid only once within each period (*e.g.*, a day).

We denote the number of entrants and sellers in period t as $\mathcal{C}_t^N$ and $\mathcal{S}_t^N$ respectively.[8] 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 $k \in \{0,1,...\}$ buyers from the buyer queue where k is a realization of random variable K that is distributed according to probability mass function $\pi(K;\boldsymbol{\lambda})$. For now, we impose no functional form on π, meaning the parameter vector $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\} \in \mathbb{R}^\infty$ is left unrestricted.

We set the mean of K to be equal to $E[K] = \mathcal{C}_t^N/\mathcal{S}_t^N$, and this condition implicitly determines $\boldsymbol{\lambda}$. This leads us to refer to $\boldsymbol{\lambda}$ as the *market tightness* parameter since it is determined by the buyer-seller ratio. Intuitively, 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 buyers*. Unmatched buyers proceed to the next period without transacting.

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

(1) Bidders are matched to an auction almost surely as the number of players grows

(2) π is continuous in $\boldsymbol{\omega}$, F_V, and G_R

[8] Since all sellers enter, $\mathcal{S}_t^N = S_t^N$. We employ this notation for expositional simplicity.

(3) *A local limit theorem applies, meaning that for the sequence $(K_1, K_2, ...)$ with $Z_N = \sum_{i=1}^{N} K_i$ and ψ denoting the density function of the normal distribution we have*

$$\sqrt{NVar[K]}Pr\{Z_N = k\} \to \psi\left(\frac{k - NE[K]}{\sqrt{NVar[K]}}\right) \textit{ uniformly over } k \in \mathbb{Z} \tag{1}$$

The most novel assumption is item (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 many distributions of common interest including, for example, the generalized Poisson distribution used in our estimator.[9] 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 about the total number of competitors from the perspective of a bidder in the auction are not the same as beliefs from the perspective of an outside observer (*e.g.*, a seller or the platform designer). From the perspective of a bidder matched into an auction, let M be a random variable representing the number of competitors she faces and $\pi_M(M; \boldsymbol{\lambda})$ denote its 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, the bidder's beliefs about M are

$$\pi_M(m; \boldsymbol{\lambda}) = Pr\left[m \text{ opponents}|\lambda\right] = \pi(m+1; \boldsymbol{\lambda})\frac{(m+1)}{\mathrm{E}[K]}$$

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.

2.2. **Equilibrium Bids.** In this section we discuss the structure of the equilibrium under the assumption of a second-price auction (SPA) rule since we can provide closed form solutions for some equilibrium quantities. As we show in Section 4.1.3, our general insights apply straightforwardly to other pricing mechanisms as well.

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

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

We use the notation $x^N(b, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N) = 1$ (0) to denote the random event that a buyer wins (loses) an auction with a bid of b, and $p^N(b, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N)$ denotes the random transfer from the buyer to the seller/eBay conditional on a bid of b.[10] To simplify notation we also define

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

Note that ρ represents expected transfers that are not conditional on sale. That is, each entering bidder has an *ex-ante* expectation to pay $\rho^N(b)$ 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 often also suppress the bid argument and assume the agent is following the equilibrium (or candidate equilibrium) strategy. We superscript the expectation operator to emphasize that we are referring to the N-agent game.

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

$$\mathcal{V}^N(v_i, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N | \sigma) = \chi^N v_i - \rho^N - \kappa + (1 - \chi^N)\delta E_t^N \left[\mathcal{V}^N(v_i, \boldsymbol{\omega}_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma) \right]$$

For a buyer that chooses *Out* we have

$$\mathcal{V}^N(v_i, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N | \sigma) = \delta E_t^N \left[\mathcal{V}^N(v_i, \boldsymbol{\omega}_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma) \right]$$

We use the notation $\mathcal{V}^N(v_i, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N | \sigma_i', \sigma_{-i})$ when buyer i uses strategy σ' and all other agents follow σ.

When facing an SPA mechanism, it is an equilibrium in weakly un-dominated strategies for a bidder to bid his value for the good minus the opportunity cost of winning. In the static, one-shot setting, the 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

$$\mathcal{O}(v_i, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N) = v_i - \delta E_t^N \left[\mathcal{V}^N(v_i, \boldsymbol{\omega}_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N | \sigma) \right]$$

We use the following definition of an equilibrium in our finite games.[11]

---

[10]For example, $p^N(b, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N) = 0$ if the buyer does not win the auction.

[11]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

Definition 2.2. The strategy vector $\sigma = (\theta, \mathcal{O})$ and the initial state $\boldsymbol{\omega}_0^N \in \Omega_N$ 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

$$\text{For all } \sigma_i' \in \Sigma,\ \mathcal{V}^N(v_i, \boldsymbol{\omega}_0^N, F_{V,0}^N, G_{R,0}^N|\sigma) + \varepsilon \geq \mathcal{V}^N(v_i, \boldsymbol{\omega}_0^N, F_{V,0}^N, G_{R,0}^N|\sigma_i', \sigma_{-i})$$

We now show that there exists a 0-BNE for our large finite model if there exists an equilibrium for the static version of our auction model (*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 to 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 is her private value minus the opportunity cost of winning. Proposition 2.3 makes it possible to apply our identification strategy to non-SPA pricing rules.

Proposition 2.3. *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*

$$\theta(v, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N)) = \tilde{\theta}(v - \delta E_t^N\left[\mathcal{V}^N(v, \boldsymbol{\omega}_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N|\sigma)\right], \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N)$$

$$\mathcal{O}(v, \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N)) = \tilde{\mathcal{O}}(v - \delta E_t^N\left[\mathcal{V}^N(v, \boldsymbol{\omega}_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N|\sigma)\right], \boldsymbol{\omega}_t^N, F_{V,t}^N, G_{R,t}^N)$$

3. CONTINUUM MODEL

In the finite model, the agents condition their actions on the state of the economy (*i.e.*, $\boldsymbol{\omega}_t$, $F_{V,t}^N$, and $G_{R,t}^N$). The set of values that these variables can assume grows exponentially with N. The finite model suffers a curse of dimensionality in the sense that computing an equilibrium requires describing a strategy that is a best response to each element of the exponentially growing set of possible states. This makes the the finite model intractable to solve for even moderately large numbers of agents.

In the continuum model that we introduce in this section, the evolution of the state of the economy is deterministic, but must be consistent with the equilibrium strategy. What makes the estimation and simulation of the continuum model tractable is that the equilibrium strategies need only condition on the deterministic path of the state of the economy. After establishing existence of a stationary equilibrium of our limit

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. As we discuss in more detail below, the thrust of our analysis is to compute stationary equilibria of the limit game. A stationary strategy can be an ε-BNE even though perfection is not even approximately satisfied at histories of the finite game in which the market aggregates differ significantly from the stationary state.

model, we show that it is an ε-BNE of the large finite model of interest and that we can take $\varepsilon \to 0$ as $N \to \infty$. As a result, we use the terms "continuum model" and "limit model" interchangeably. Associated proffs are lengthy and therefore relegated to an online technical appendix.

3.1. **The Continuum Model.** Our continuum model features positive measures of buyers and sellers that are individually of measure 0. The measure of buyers in period $t = 0$ is C_0^∞, and the type distribution of the buyers at $t = 0$ is $F_{V,0}$. 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 entrant buyers is equal to T_V. In each period there is a measure S_t^∞ of sellers with reserve prices distributed exactly as G_R. We let $\boldsymbol{\omega}_t = (C_t^\infty, S_t^\infty)$.

If a measure $\mathcal{C}_t^\infty$ of buyers choose to enter the auction market, the buyers are randomly assigned to auctions with each auction receiving k bidders with probability $\pi(k, \boldsymbol{\lambda})$ satisfying the regularity conditions mentioned in the previous section and $\mathrm{E}[K] = C_t^\infty / S_t^\infty$. Given the number of bidders assigned to each seller's auction, the allocation and price setting procedure within each auction is executed as in the finite model. 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. We use similar notation to describe the limit model, although we drop the superscript N when referring to objects pertinent to the continuum model. We use the notation

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

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. Again we supress the aggregate variables and the specific bid when confusion will not result.

The value function in the continuum game given a (symmetric) equilibrium strategy vector $\sigma = (\theta, \mathcal{O})$ for a buyer that chooses *Enter* is

$$\mathcal{V}(v_i, \boldsymbol{\omega}_t, F_{V,t}, G_R|\sigma) = \chi v_i - \rho - \kappa + (1-\chi)\delta \mathcal{V}(v_i, \boldsymbol{\omega}_{t+1}, F_{V,t+1}, G_R|\sigma),$$

and for a buyer that chooses *Out* we have

$$\mathcal{V}(v_i, \boldsymbol{\omega}_t, F_{V,t}, G_R|\sigma) = \delta \mathcal{V}(v_i, \boldsymbol{\omega}_{t+1}, F_{V,t+1}, G_R|\sigma).$$

Note the absence of the expectation operators in the limit model. This is because the evolution of market aggregates is deterministic once a strategy and $(\omega_0, F_{V,0}, G_R)$ have been fixed.[12] Moreover, the evolution of market aggregates is unaffected by the actions of any single measure 0 agent.

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 economic aggregates are constant across time and the agent actions are optimal with respect to the fixed aggregate state the agents face. θ refers to the stationary entry strategy of the buyers, and β refers to a stationary bidding strategy.

Definition 3.1. The strategy vector $\sigma = (\theta, \beta)$ and the state $\omega \in \Omega$ and $F_V \in \Delta([0,1])$ is a Stationary Competitive Equilibrium (SCE) if for all values v we have

(1) For all $\sigma_i' \in \Sigma$,

$$\mathcal{V}(v_i, \boldsymbol{\omega}, F_V, G_R|\sigma) \geq \mathcal{V}(v_i, \boldsymbol{\omega}, F_V, G_R|\sigma_i', \sigma_{-i})$$

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

$$\chi v_i - \rho - \kappa + (1-\chi)\delta\mathcal{V}(v_i, \boldsymbol{\omega}_{t+1}, F_V, G_R|\sigma) \geq 0$$

(3) $\boldsymbol{\omega} = \boldsymbol{\omega}_t$ and $F_V = F_{V,t+1}$ are consistent with the laws of motion of the game.

In a stationary equilibrium, the agents bid the same amount in each period, meaning that the bidding function can be written $\beta : [0,1] \to [0,1]$. In equilibrium, the entry decision must take the form

$$\theta(v_i, \boldsymbol{\omega}, F_V, G_R) = \textit{Enter} \text{ if and only if } \chi\left(v_i - \delta\mathcal{V}^{\infty}(v_i, \boldsymbol{\omega}, F_V, G_R|\sigma)\right) - \rho \geq \kappa$$

Any buyer that is indifferent between entering and staying out must have a continuation value of 0 since, due to 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 $e(\boldsymbol{\omega}, F_V, G_R) = \inf_v\{v : \chi v - \rho \geq \kappa\}$ and

$$\theta(v_i, \boldsymbol{\omega}, F_V, G_R) = \textit{Enter} \text{ if and only if } v \geq e(\boldsymbol{\omega}, F_V, G_R) \tag{2}$$

Once again, under the SCE with a second-price spot-market rule, buyers submit bids equal to the opportunity cost of re-entering the market next period:

$$\beta(v_i) = v_i - \delta V(v_i, \boldsymbol{\omega}, F_V, G_R|\sigma) \tag{3}$$

[12]The uncertainty regarding the events that will take place in the current period are accounted for through our use of the expected allocation (χ) and expected transfer (ρ).

In the interest of generality, we extend our framework to account for a variety of payment rules. For our large-market approximation results to hold for platforms with non-second-price spot-market mechanisms, we require the following assumptions on the equilibrium strategies:

Assumption 3.2. *The best responses by the buyers in the continuum game are continuous in the sup-norm with respect to the parameter $\boldsymbol{\lambda}$ 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 (ω, g_B, G_R) and any $v > v'$ we have

$$b(v) - b(v') \in \left[\varphi(v - v'), \frac{v - v'}{\varphi}\right] \tag{4}$$

For the duration of this paper, we take assumption 3.2 as given. Given Theorem 2.3, we could have equivalently required continuity of the bidding equilibrium of the static game ($\delta = 0$).[13] In summary we can describe any stationary equilibrium as a vector of strategies, $\sigma = (e, \beta)$, and stationary aggregate states ,$(\boldsymbol{\omega}, F_V, G_R)$. Our next result shows that there exists a stationary equilibrium of the continuum model.

Proposition 3.3. *A stationary competitive equilibrium exists in the continuum model, 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) proving we can limit consideration to a compact strategy space and (2) ruling out a number of utility discontinuities that naturally arise in auction markets. Once we handle these issues, our proof relies on a traditional fixed point argument.

3.2. **Approximating the Large Finite Model.** Our goal is to prove that the limit model approximates the large finite model. The foundation of our proofs is a mean field result that proves that the evolution of the limit game economy and the economy of a finite model 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, and proving discontinuities do not arise is challenging.

We must also demonstrate that the expected buyer utility in the large finite game and the limit game are approximately the same, which is challenging in our model. Since the

[13]In the SPA we have $\tilde{\beta}(v) = v$, which clearly satisfies assumption 3.2. Since the value function must be strictly increasing with a slope less than 1, we know that $v - \delta\mathcal{V}(v)$ satisfies Equation (4). Therefore the SPA satisfies assumption 3.2.

within-period matching process samples without replacement from a finite set of buyers, there is correlation across auction outcomes in the finite game that is not present in the limit game. Moreover, there is also a positive probability that a positive mass of buyers is unmatched in the finite game. We show that these problems vanish as the size of the market increases.

We can translate these insights into our approximation result, which proves that any exact equilibrium strategy of the limit game is an $\varepsilon-$BNE of the finite game with sufficiently many players.

Proposition 3.4. *Consider an SCE $(\sigma, \boldsymbol{\omega}, F_V, G_R)$ and assume $e(\boldsymbol{\omega}, F_V, G_R) < 1$. For any $\varepsilon > 0$ we can choose N^* and $\eta > 0$ such that for all $N > N^*$, σ is an $\varepsilon - BNE$ strategy if $\left(\omega_0^N, F_{V,0}^N, G_{R,0}^N\right)$ satisfies*

$$\left\|\boldsymbol{\omega}_0^N - \boldsymbol{\omega}\right\| + \left\|F_{V,0}^N - F_V\right\| + \left\|G_{R,0}^N - G_R\right\| < \eta \tag{5}$$

This result is significant because it establishes that a bidder's value function within the finite game is approximated by the limiting value function, which is the justification for why we use the limit model as the basis of our empirical framework. Proposition 3.4 may be seen as providing an approximation to the actual equilibrium being played within the data-generating process, but it admits an alternative interpretation of a behavioral strategy as well. If one assumes that agents are subject to small computation 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 $N * C^\infty$ and $N * S^\infty$, then $\left(\boldsymbol{\omega}_0^N, F_{V,0}^N, G_{R,0}^N\right) \to (\boldsymbol{\omega}, F_V, G_R)$ almost surely as $N \to \infty$. In other words, Equation (5) above becomes increasingly likely to hold as $N \to \infty$. Moving forward, we assume that our data-generating process is of sufficient size for the limiting model to provide an accurate approximation to bidders' equilibrium payoffs.

4. AN EMPIRICAL MODEL OF DYNAMIC PLATFORM MARKET BIDDING

We now shift focus to developing a structural model based on the assumption that bidders within the finite (but large) marketplace of eBay behave according to the approximately optimal SCE described above. Letting L denote sample size (where an auction is the unit of observation), the observables, $\{\tilde{k}_l, r_l, y_l\}_{l=1}^{L}$, are assumed to 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 of identification, we leave the bidder arrival process $\pi(\cdot; \boldsymbol{\lambda})$ nonparametric and the market tightness parameter vector $\boldsymbol{\lambda}$ is allowed to be infinite-dimensional with $\lambda_k = \Pr[K = k],\ k = 0, 1, 2, \ldots$. For convenience we drop the parameter argument in the PMF $\pi(\cdot)$ unless context requires specificity. Following the previous section, we assume measure μ of new buyers exogenously flow into the market each period, and we also assume that the various type distributions—T_V, F_V, G_R—are maintained in steady state, where the following identity ensures a stationary market:

$$\mu t_V(v) = \chi(\beta(v)) f_V(v) \frac{\mathcal{C}^\infty}{\mathcal{S}^\infty} = \chi(\beta(v)) f_V(v) \mathrm{E}[K]. \tag{6}$$

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 b, and $\mathcal{C}^\infty/\mathcal{S}^\infty$ is the market-wide buyer-seller ratio. The left-hand side of Equation (6) 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 after transacting.

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 $\pi(\cdot)$. Prior to bidding, 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 $[\underline{v}, \overline{v}]$, with $\underline{v} > 0$.[14] 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 (2) determining the relevant cutoff. However, 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 T_V and F_V are steady-state distributions for buyer types who choose *Enter*. This is possible because in an SCE of the limit model, if 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. Therefore, $\underline{v}$ is interpreted as the infra-marginal type who is just indifferent to entering, and our counterfactual investigation will focus only on scenarios involving subsets of current market entrants. We normalize payoffs of non-entrants to zero, leading to the following which we refer to as the "zero surplus condition:"

Assumption 4.1. $\mathcal{V}(\underline{v}) = 0$.

Bidder 1 views the bids of her opponents as a random variable $B = \beta(V) \sim G_B(B) = F_V\left[\beta^{-1}(B)\right]$ with support $[\underline{b}, \overline{b}]$. Let B_M denote the maximal bid among all of bidder 1's

[14] Nothing 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 interval of real numbers.

opponents with distribution defined by $G_{B_M}(B_M) = \sum_{m=1}^{\infty} \frac{\pi_M(m)}{1-\pi_M(0)} G_B(B_M)^m$. In order to win, player 1's bid must exceed the realized value of a random variable

$$Z \equiv \begin{cases} R & \text{if } M = 0, \\ \max\{R, B_M\} & \text{otherwise.} \end{cases}$$ [15]

Note then that the distribution of Z is the same as the win probability function, or

$$\chi(b) \equiv \mathrm{E}[x_B(b)] = G_R(b) \sum_{m=0}^{\infty} \pi_M(m) G_B(b)^m \tag{7}$$

4.1. **Model Identification.** We first establish nonparametric identification results in Sections 4.1.1–4.1.2 for a baseline case with second-price spot-market auctions. In this setting a buyer's tendency for demand shading derives entirely from dynamic incentives. In Section 4.1.3 we extend our identification result to the case where the spot-market game is non-second price, which creates additional 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.

4.1.1. *Baseline Model: Second-Price, Sealed-Bid Spot-Market Auctions.* A second-price spot-market mechanism implies an expected payment function of the form

$$\begin{aligned} \rho(b) \equiv \mathrm{E}[p_B(b)] = &\pi_M(0) G_R(b) \mathrm{E}[R | R \leq b] \\ &+ [1 - \pi_M(0)] \int_{\underline{b}}^{b} t \left[g_R(t) G_{B_M}(t) + G_R(t) g_{B_M}(t) \right] dt. \end{aligned} \tag{8}$$

Since the market is in steady-state, we can simplify notation from the previous section and express the Bellman equation and dynamic bidding strategy as

$$\mathcal{V}(v) = \max_{b \in \mathbb{R}_+} \left\{ \chi(b) v - \rho(b) - \kappa + [1 - \chi(b)] \, \delta \mathcal{V}(v) \right\} \text{ and} \tag{9}$$

$$\beta(v) = v - \delta \mathcal{V}(v). \tag{10}$$

In other words, bidders in the dynamic, second-price platform market shade their demand according to the option value of returning to the market the next period in the event of a loss. The demand shading factor, given by bidder 1's continuation value,

[15]Recall that within the model bidders formulate their bids *before* being matched to an auction and observing the realization of R, which is independent of B. This assumption is motivated by the idea that intertemporal incentives play a predominant role in determining bidders' payoffs, and therefore they feel no need to update strategies according to the finer details of any one given auction. One implication of the model provides a means for a partial test of this assumption. If R and B are independent random variables, then the highest losing bid and the reserve price should be uncorrelated. We find in our data that the correlation coefficient between the highest losing bid and the reserve price is small (-0.015) and statistically indistinguishable from zero.

$\delta\mathcal{V}(v)$, is uniquely characterized by three things: the per-period entry cost, κ, the distribution of bids, $G_B(b)$, and the market tightness parameters $\boldsymbol{\lambda}$ that determine the overall ratio of buyers and sellers. Thus, mapping bids into private values requires first identifying these three 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 $\boldsymbol{\lambda}$, κ, F_V, and T_V.

Note that (10) implies $\mathcal{V}(v) = \frac{v-\beta(v)}{\delta}$. By substituting this expression into Equation (9), and using the shorthand $b^* = \beta(v)$, we can rearrange terms to get

$$v = b^* \frac{1-\delta\left(1-\chi(b^*)\right)}{1-\delta} - \frac{\delta}{1-\delta}\Big(\rho(b^*)+\kappa\Big) \equiv \beta^{-1}(b^*). \tag{11}$$

The expression for the inverse bidding function above demonstrates that if the econometrician can identify $\boldsymbol{\lambda}$, 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 any bid b as a best response to prevailing market conditions.

4.1.2. *Identification of the Bid Distribution and Bidder Arrival Process.* One challenge to our 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 intend to compete, some may find that their planned bid was surpassed before they had a chance to submit it to the server. These bidders will never be visible to the econometrician, even though they were matched to the auction and competing for the sale.

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.[16] For a given auction with k total matched bidders, this filter process first randomly assigns each bidder an index $\{1, 2, \dots, k\}$ that 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

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

some bidders will not appear to have participated, even though they had participated by watching the item with intent to bid. This view of intra-auction dynamics assumes that the ordering of bidders' submission times is random and not a function of bidder characteristics.

By explicitly modeling this source of sample selection, one can nonparametrically identify $\boldsymbol{\lambda}$ from the observed lower bounds $\tilde{k}$. 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 $\boldsymbol{\lambda}$. Moreover, since a bidder's visibility to the econometrician only depends on whether her bid exceeds the second-highest preceeding bid, the researcher can easily simulate the filter process without knowing G_B or F_V to compute conditional probabilities $\Pr[\tilde{k}|k]$ for various $(\tilde{k}, k)$ pairs.[17] 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 $\boldsymbol{\lambda}$: $\tilde{\pi}(\tilde{k}) = \sum_{k=\tilde{k}}^{\infty} P_0(\tilde{k}, k)\pi(k; \boldsymbol{\lambda})$.

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 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 witholds 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 as 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. 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,[18] and Nature then compiles a reported list of bidders for the econometrician in two steps. First, she visits each bidder and dismisses anyone whose strategic bid does not meet the reserve price r. Second, Nature assigns the remaining set of $k' \leq k$ bidders random indices $i \in \{1, 2, \ldots, k'\}$, and then executes

[17] Hickman 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 $\mathrm{E}[K] \leq 40$.

[18] Recall that the bidders formulated their bids before being assigned to a particular auction.

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 $\binom{k}{j} G_B(r)^j \left[1 - G_B(r)\right]^{k-j}$. 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:

$$\Pr[\tilde{K} = \tilde{k} | K = k, r] = \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} G_B(r)^j \left[1 - G_B(r)\right]^{k-j} P_0(\tilde{k}, k-j) \tag{12}$$

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 accounts for the standard filter process running its course with the surviving set of bidders. Equation (12) allows us to characterize the distribution of observed $\tilde{K}$ conditional on the observable reserve price r, as

$$\tilde{\pi}(\tilde{k}|r) = \sum_{k=\tilde{k}}^{\infty} \Pr[\tilde{k}|k, r] \pi(k; \boldsymbol{\lambda}). \tag{13}$$

Estimation of $\boldsymbol{\lambda}$ can no longer be separated from G_B because Equation (13) above 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 4.2. *For a given discount factor δ, the market tightness parameters $\boldsymbol{\lambda}$, 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

$$H(b) = \sum_{k=2}^{\infty} \frac{\pi(k; \boldsymbol{\lambda})}{1 - \pi(0; \boldsymbol{\lambda}) - \pi(1; \boldsymbol{\lambda})} \left(G_B(b)^k + k G_B(b)^{k-1} \left[1 - G_B(b)\right] \right). \tag{14}$$

$H(b)$ 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 $\varphi\left(H(b); \boldsymbol{\lambda}\right) = G_B(b)$ denote the inverse of (14), then it follows that, holding $\boldsymbol{\lambda}$ 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 $\boldsymbol{\lambda}$ and G_B

are identified from the observables and from equations (13) and (14). 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 (7) and (8) above. To identify the participation cost, combine Equation (9) with the zero surplus condition (4.1) to find the following relation

$$\chi(\underline{v})\underline{v} - \rho(\underline{v}) = \kappa \tag{15}$$

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 (11) 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:

$$\begin{aligned} \mu &= \int_{\underline{v}}^{\overline{v}} \chi\left[\beta(v)\right] f_V(v) dv \\ &= \left[1 - \pi(0)\right] G_R(\underline{b}) + \int_{\underline{b}}^{\overline{b}} g_R(r) \left(\sum_{k=1}^{\infty} \pi(k) \left[1 - G_B(r)^k\right] \right) dr \end{aligned} \tag{16}$$

Finally, once μ is known T_V is identified through Equation (6). □

4.1.3. *Model Identification Under Alternative Spot Market Mechanisms.* We now extend our identification result to cover platform markets that use alternative spot-market pricing mechanisms. Alternative spot-market mechanisms in which the winner's bid directly influences the current-period sale price will produce incentives for demand shading above and beyond the dynamic demand shading incentives described above. Proposition 2.3 combined with identification arguments from the empirical literature on static auctions allow us to disentangle these incentives and identify the bidders' underlying valuations.

We shift notation slightly and use $\rho(b)$ to denote the expected payment under the prevailing spot market mechanism, whatever it may be. We find it useful to refer to a bidder's private value minus her opportunity cost as her *dynamic value*, denoted $\tilde{v}_v \equiv v - \delta\mathcal{V}(v)$.. Proposition 2.3 implies that we can re-cast the bidder's decision problem as choosing a functional $\tilde{\beta} : \tilde{V}_V \to \mathbb{R}_+$ to optimize an alternative but equivalent objective

$$\tilde{\beta}(\tilde{v}_v) = \arg\max_b \left\{ \chi(b)\tilde{v}_v - \rho(b) \right\}, \tag{17}$$

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 $\tilde{v}_v$. 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 are also identified. To see why, note that by plugging the optimizer $\tilde{\beta}$ into Equation (9) 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

$$v = \tilde{v}_v \left(\frac{1 - \delta \left[1 - \chi\left(b^*\right)\right]}{1 - \delta} \right) - \frac{\delta}{1 - \delta} \Big(\rho\left(b^*\right) + \kappa \Big) = \beta^{-1}(b^*). \tag{18}$$

In the case of a second-price spot market, Equation (18) reduces to Equation (11) above.

Proposition 4.3. *For a given discount factor δ, the market tightness parameters $\boldsymbol{\lambda}$, 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 (17) 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 (17) 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 $\boldsymbol{\lambda}$, G_B, and G_R is the same as in Proposition 4.2. For case (1), assuming that $\chi(b)$ and $\rho(b)$ can be expressed as a function of observable objects (including $\boldsymbol{\lambda}$, G_B, and G_R), equations (15) and (17) 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 $\boldsymbol{\lambda}$, 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 (18) 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 (6) and (16) identify μ and T_V similarly as before. □

Proposition 4.3 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 in which the pricing and allocation rules can be expressed in terms of $\boldsymbol{\lambda}$, G_R, and G_B. 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 $\pi(\cdot;\boldsymbol{\lambda})$. The ability to incorporate established identification strategies for static auctions will be useful as we develop an estimator for eBay data, which uses a pricing rule that is a non-standard combination of both first-price and second-price rules.

4.2. **A Two-Stage, Semi-Parametric Estimator.** Thus far in our discussion we have left the bidder arrival process $\pi(k;\boldsymbol{\lambda})$ 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 assume K follows a generalized Poisson distribution with PMF

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

due to Consul and Jain [1973]. One advantage to our approach described below is that it is easily adaptable to more complex functional forms if the researcher is concerned about mis-specification problems. A fully nonparametric estimator involves additional complications beyond the scope of this work; see online empirical appendix for a breif discussion. The first two moments of the generalized Poisson distribution are $\mathrm{E}[K] = \lambda_1/(1-\lambda_2)$ and $\mathrm{Var}[K] = \mathrm{E}[K]/(1-\lambda_2)^2$. Intuitively, the parameter λ_2 regulates the dispersion of the random variable K. The generalized Poisson reduces to a regular Poisson distribution when $\lambda_2 = 0$, but exhibits fatter tails when $\lambda_2 > 0$ and thinner tails when $\lambda_2 < 0$. Developing an estimator based on finite-dimensional $\boldsymbol{\lambda}$ avoids significant complications that we discuss briefly below, but which are beyond the scope of this work.

Recall that bidder 1's beliefs about the number of her opponents, M, follows

$$\pi_M(m,\boldsymbol{\lambda}) = \pi(m+1;\boldsymbol{\lambda})(m+1)\frac{(1-\lambda_2)}{\lambda_1}$$

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 $\boldsymbol{\lambda}$ 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.[19] B-splines provide a remarkable combination of flexibility and numerical convenience that is ideally suited to our application.

Let $\mathbf{n}_b = \{n_{b1} < n_{b2} < \cdots < n_{b,I_b+1}\}$ be a set of knots on bid domain $[\hat{\underline{b}}, \hat{\overline{b}}] = [\min_l\{y_l\}, \max_l\{y_l\}]$ that create a partition of I_b subintervals. This need not be a uniform partition, but we do require that $n_{b1} = \hat{\underline{b}}$ and $n_{b,I_b} = \hat{\overline{b}}$ 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 $\mathcal{F}_{bi} : [\hat{\underline{b}}, \hat{\overline{b}}] \to \mathbb{R},\ i = 1, \ldots, I_b + 3$ that give us our parameterization of the bid distribution:

$$\hat{G}_B(b; \boldsymbol{\alpha}_b) = \sum_{i=1}^{I_b+3} \alpha_{b,i} \mathcal{F}_{bi}(b).^{20}$$

We also follow this approach for estimating G_R and F_V. Let $\mathbf{n}_r = \{n_{r1} < n_{r2} < \cdots < n_{r,I_r+1}\}$ and $\mathbf{n}_v = \{n_{v1} < n_{v2} < \cdots < n_{v,I_v+1}\}$ 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 $[\underline{r}, \hat{\overline{r}}] = [0.99, \max_l\{r_l\}]$ and the latter spans $[\hat{\underline{v}}, \hat{\overline{v}}]$, with the bounds to be estimated. These knot vectors determine our other basis functions $\mathcal{F}_{ri} : [\underline{r}, \hat{\overline{r}}] \to \mathbb{R},\ i = 1, \ldots, I_r + 3$ and $\mathcal{F}_{vi} : [\hat{\underline{v}}, \hat{\overline{v}}] \to \mathbb{R},\ i = 1, \ldots, I_v + 3$ which in turn render our parameterizations $\hat{G}_R(r; \boldsymbol{\alpha}_r) = \sum_{i=1}^{I_r+3} \alpha_{ri} \mathcal{F}_{ri}(r)$ and $\hat{F}_V(v; \boldsymbol{\alpha}_v) = \sum_{i=1}^{I_v+3} \alpha_{vi} \mathcal{F}_{vi}(v)$.

Following our identification argument, we separate estimation into two stages. In the first stage we flexibly estimate $\boldsymbol{\lambda}$, 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 1 and 2 differ in that Stage 1 is an estimation step, but Stage 2 is a purely computational step based on the outputs from Stage 1.

4.2.1. *Stage 1: $\boldsymbol{\lambda}$, 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

[19] Unlike 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.

[20] A 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.

model-generated conditional PMF of $\tilde{K}$ given r as

$$\tilde{\pi}(\tilde{k}|r;\boldsymbol{\lambda},\boldsymbol{\alpha}_b) = \sum_{k=\tilde{k}}^{\overline{K}} \left\{ \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} \hat{G}_B(r;\boldsymbol{\alpha}_b)^j \left[1-\hat{G}_B(r;\boldsymbol{\alpha}_b)\right]^{k-j} P_0(\tilde{k},k-j) \right\} \pi(k;\boldsymbol{\lambda}).$$

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

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

where $\mathbb{1}(\cdot)$ is an indicator function, $\mathcal{K}$ is a boundary-corrected kernel function, and h_R is an appropriately chosen bandwidth.[21] Finally, we define the model-generated highest loser bid distribution as

$$H(b;\boldsymbol{\lambda},\boldsymbol{\alpha}_b) = \sum_{k=2}^{\infty} \frac{\pi(k;\boldsymbol{\lambda})\left(G_B(b;\boldsymbol{\alpha}_b)^k + kG_B(b;\boldsymbol{\alpha}_b)^{k-1}\left[1-G_B(b;\boldsymbol{\alpha}_b)\right]\right)}{1-\pi(0;\boldsymbol{\lambda})-\pi(1;\boldsymbol{\lambda})},$$

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

$$\begin{aligned}(\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b) = &\underset{(\boldsymbol{\lambda},\boldsymbol{\alpha}_b)\in\mathbb{R}^{I_b+5}}{\arg\min} \sum_{l=1}^{L} \left\{ \left[\tilde{\pi}(\tilde{k}_l|r_l;\boldsymbol{\lambda},\boldsymbol{\alpha}_b) - \hat{\tilde{\pi}}(\tilde{k}_l|r_l)\right]^2 + \left[H(y_l;\boldsymbol{\lambda},\boldsymbol{\alpha}_b)-\hat{H}(y_l)\right]^2 \right\} \\ &\text{subject to} \qquad \alpha_{b1}=0,\ \ \alpha_{b,I_b+3}=1, \\ &\qquad\qquad\qquad \alpha_{b,i} \leq \alpha_{b,i+1},\ i=1,\ldots,I_b+2. \end{aligned} \tag{19}$$

In words, the estimate $(\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_G)$ is chosen to make the model-generated conditional distribution of $\tilde{K}$ match its empirical analog as closely possible.[22] The constraints on the empirical objective function enforce boundary conditions and monotonicity of our parameterization for $\hat{G}_B$.[23]

[21] The 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.

[22] hello world

[23] One 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), exactly one basis function is nonzero at the lower boundary, $\mathcal{F}_{b1}(\underline{b}) = 1$, and exactly one basis function is nonzero at the upper boundary, $\mathcal{F}_{b,I_b+3}(\underline{b}) = 1$. 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 $\hat{G}_B(b;\boldsymbol{\alpha}_b)$ 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.

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

$$\begin{aligned} \hat{\boldsymbol{\alpha}}_r = \underset{\boldsymbol{\alpha}_r \in \mathbb{R}^{I_r+3}}{\arg\min} & \sum_{l=1}^{L} \left\{ \left[\hat{G}_R(r_l; \boldsymbol{\alpha}_r) - \ddot{G}_R(r_l) \right]^2 \right\} \\ \text{subject to} \quad & \alpha_{r1} = \ddot{G}_R(\underline{r}), \ \ \alpha_{r,I_r+3} = 1, \\ & \alpha_{ri} \leq \alpha_{r,i+1}, \ i = 1, \ldots, I_r + 2, \end{aligned} \tag{20}$$

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

4.2.2. *Stage 2:* Having these estimates in hand, we are able to directly re-construct the remaining structural primitives. Some Stage 2 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.

Before moving on, a word on spot market mechanisms is in order. Prevailing wisdom in empirical work has often held 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 $\Delta > 0$. If a bidder wins, then the price will be set equal to $Z + \Delta$. However, a complication arises when the winning bid and Z are within Δ of each other, as this would involve a price $Z + \Delta$ that exceeds the winner's bid. In that case, the price is set equal to the winner's bid as in a first-price mechanism. Thus, eBay's pricing rule follows $p(b) = min\{Z + \Delta, b\}$.

Hickman [2010] proved the 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. He also showed that 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

$$v = b + \frac{G_{B_M}(b) - G_{B_M}\left[\tau(b)\right]}{g_{B_M}(b)}, \quad \tau(b) = \begin{cases} \underline{b} & \text{if } b \leq \underline{b} + \Delta \\ b - \Delta & \text{otherwise,} \end{cases} \tag{21}$$

where $\tau(b)$ is a threshold function determining the point below one's own bid that triggers a first-price outcome.

Proposition 4.3 enables us to adapt Equation (21) 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 bidder type $\underline{b}$ 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

$$\hat{\tilde{\beta}}^{-1}(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)=\hat{\tilde{v}}_b=b+\frac{\hat{G}_Z(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)-\hat{G}_Z\left[\tau(b);\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r\right]}{\hat{g}_Z(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)}. \tag{22}$$

Using Stage 1 estimates we can construct the allocation rule and the distribution of Z:

$$\begin{aligned}\chi(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)=&\hat{G}_R(b;\hat{\boldsymbol{\alpha}}_r)\sum_{m=0}^{\infty}\pi_M(m;\hat{\boldsymbol{\lambda}})\hat{G}_B(b;\hat{\boldsymbol{\alpha}}_b)^m\\ =&\hat{G}_Z(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r),\ b\geq 0,\ \text{ and}\end{aligned} \tag{23}$$

Equation (23) is a straightforward adaptation of (7), and 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:

$$\begin{aligned}\rho(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)=\underline{r}G_R(\underline{r};\hat{\boldsymbol{\alpha}}_r)+&\int_{\underline{r}}^{\tau(b)}(t+\Delta)\hat{g}_Z(t;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)dt\\ &+b\left(\hat{G}_Z[b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r]-\hat{G}_Z[\tau(b);\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r]\right).\end{aligned} \tag{24}$$

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.

At this point, we can recover the per-period entry cost as

$$\hat{\kappa}=\chi(\underline{\hat{v}_b};\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)\underline{\hat{v}_b}-\rho(\underline{\hat{v}_b};\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r), \tag{25}$$

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

$$\hat{v}=\hat{\beta}^{-1}\left(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right)=\hat{\tilde{v}}_v\frac{1-\delta\left[1-\chi(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)\right]}{1-\delta}-\frac{\delta\left(\rho(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)+\hat{\kappa}_B\right)}{1-\delta} \tag{26}$$

$$\hat{\mathcal{V}}\left(v;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right)=\frac{\hat{v}-\hat{\tilde{v}}_b.}{\delta} \tag{27}$$

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, $\mathbf{b}_J=\{b_1,\ldots,b_J\}$, and a knot vector $\mathbf{n}_v$ that spans $[\underline{\hat{v}},\hat{\bar{v}}]=\left[\hat{\beta}^{-1}\left(\underline{\hat{b}};\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right),\hat{\beta}^{-1}\left(\hat{\bar{b}};\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right)\right]$. This in turn defines our basis functions $\mathcal{F}_{vi} : [\underline{\hat{v}}, \hat{\overline{v}}] \to \mathbb{R},\ i = 1, \ldots, I_v + 3$, from which we can now compute $\boldsymbol{\alpha}_v$:

$$\hat{\boldsymbol{\alpha}}_v = \underset{\boldsymbol{\alpha}_v \in \mathbb{R}^{I_v+3}}{\arg\min} \sum_{j=1}^{J} \left\{ \left[\hat{G}_B\left(b_j; \hat{\boldsymbol{\alpha}}_b\right) - \sum_{i=1}^{I_v+3} \alpha_{vi} \mathcal{F}_{vi}\left[\hat{\beta}^{-1}\left(b_j; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b, \hat{\boldsymbol{\alpha}}_r, \delta\right)\right] \right]^2 \right\}$$
$$\text{subject to} \qquad \alpha_{v1} = 0, \quad \alpha_{v,I_v+3} = 1, \qquad\qquad (28)$$
$$\alpha_{vi} \leq \alpha_{v,i+1},\ i = 1, \ldots, I_v + 2.$$

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

$$\hat{\mu} = \left[1 - \pi\left(0; \hat{\boldsymbol{\lambda}}\right)\right] G_R\left(\underline{b}; \hat{\boldsymbol{\alpha}}_r\right) + \int_{\underline{b}}^{\overline{b}} g_R\left(r; \hat{\boldsymbol{\alpha}}_r\right) \left(\sum_{k=1}^{\infty} \pi\left(k; \hat{\boldsymbol{\lambda}}\right) \left[1 - G_B\left(r; \hat{\boldsymbol{\alpha}}_b\right)^k\right] \right) dr \qquad (29)$$

$$t_V\left(v; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b; \hat{\boldsymbol{\alpha}}_r, \delta\right) = \frac{\chi\left[\beta\left(v; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b, \hat{\boldsymbol{\alpha}}_r, \delta\right); \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b, \hat{\boldsymbol{\alpha}}_r\right] f_V\left(v; \boldsymbol{\alpha}_v, \delta\right) \frac{\lambda_1}{1-\lambda_2}}{\hat{\mu}}. \qquad (30)$$

4.2.3. *Asymptotics and Standard Errors.* In an appendix we argue that our Stage 1 estimators $\hat{\boldsymbol{\lambda}}$, $\hat{\boldsymbol{\alpha}}_b$, and $\hat{\boldsymbol{\alpha}}_r$ 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 2 empirical objects are all smooth functions of Stage 1 parameters, it follows that they are also asymptotically normal, and their standard errors can be computed via the delta method. See the online empirical appendix for a detailed discussion on computation of standard errors.

5. DATA AND RESULTS

We now implement our estimator using 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 auction 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 auction listings of this item for an average of 11.25 per day. Each device is a second-generation Kindle Fire tablet (original release date: September 14, 2012) listed as new and unused or only lightly used. At the time, Amazon.com only offered one configuration of the Kindle Fire.[24]

Each Kindle Fire tablet came pre-loaded with Amazon's proprietary version of the Android-based operating system that prevents the user from accessing the full Android

[24]Each device in our sample had a 7″ screen (1024 × 600 resolution), dual-core processor (1.2GHz clock speed), 1GB RAM, 8GB internal storage, 802.11 b/g/n Wi-Fi, and ran on the Android 4.0.3 Ice Cream Sandwich operating system.

app market, which makes the Kindle Fire a poor substitute for a standard tablet (*e.g.*, Samsung Galaxy or Apple iPad) which can serve a dual role as a productivity tool or as a highly versatile consumer electronic device.[25] 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, audiobooks, music, and movies.[26] Moreover, all transactions during the sample period were covered by the eBay Money Back Guarantee to insure consumers against unscrupulous sellers.[27] These characteristics provide us with an extensive dataset for a product with no existing close substitutes and a remarkable degree of listing homogeneity. This allows us to avoid significant complications covered by other work, such as identifying unobserved heterogeneity (Krasnokutskaya [2011] or Roberts [2013], for example) or complex substitution patterns (Backus and Lewis [2016]) and instead focus on questions of allocative efficiency and market design.

5.1. **Practical Concerns.**

5.1.1. *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. The question of intra-auction dynamics is broad, complicated, and beyond the scope of this work.[28] In our case, inter-auction dynamics are the primary concern for answering our research questions on allocative efficiency and market design.

We assume our bidders formulate their strategic bids and behave as if they are competing in a simultaneous-move auction each period.[29] To deal with observed early low bids, we adopt the approach of Bajari and Hortaçsu [2003] by partitioning individual

[25] It requires specialized knowledge to un-install the proprietary operating system, and doing so is costly since it invalidates all product guarantees issued by Amazon.com.

[26] Amazon.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/Google_Play`; and `https://en.wikipedia.org/wiki/Amazon_Appstore`; information retrieved on 7/15/2016.

[27] As of 7/15/2016, details on eBay's consumer protection program were available at http://pages.ebay.com/ebay-money-back-guarantee/questions.html.

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

[29] See Figure 1 and related discussion for evidence that incentives arising from inter-auction dynamics are more important than incentives specific to the spot market itself.

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

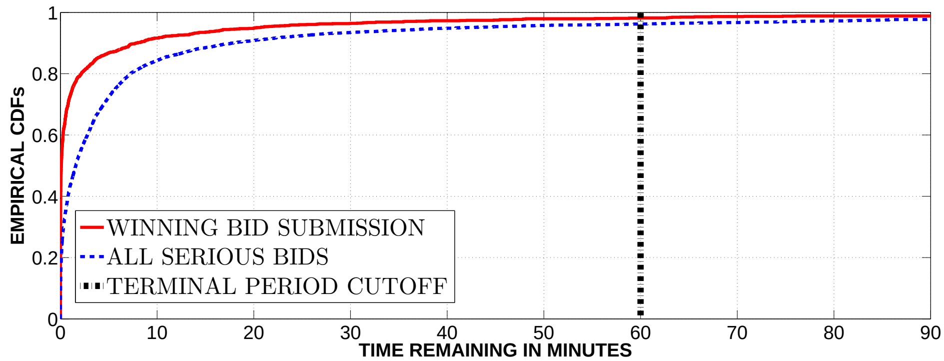

auctions into two stages. During the first phase bidders may submit cheap-talk bids that are viewed as uninformative of the other bidders' final bids and the final sale price. The second stage is treated as a sealed-bid auction.[30] The bidders optimize their planned bids and decline to update their bids within-period. 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 merely treated as 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, so it is not the case that our analysis completely ignores what occurs prior to the second stage 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

[30]While 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

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 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 parse 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.

5.1.2. *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 $\boldsymbol{\alpha}_b$ 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 1 we chose $I_b = 10$, and we partitioned the reserve price support by the quintiles of the empirical conditional distribution $\ddot{G}_R(r|R > \underline{r})$, meaning $I_r = 5$.[31] 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 1. 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.[32] 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 1 estimation, so all of the necessary building blocks to compute the final structural primitives will be unaffected. Several Stage 2 objects are also unaffected, including the static bid function $\tilde{\beta}(\cdot)$, 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 including the dynamic bid function $\beta(\cdot)$, the value function $\mathcal{V}(\cdot)$, the steady-state private value distributions for market participants $\hat{F}_V(\cdot)$ and new entrants $T_V(\cdot)$ depend on δ. 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.

[31] The conditioning is due to the mass point at the lower bound.

[32] A 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 $L \to \infty$. This is an interesting econometric question, but one which is beyond the scope of this paper.

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)

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 choice of δ. For example, instead of providing a dollar value for deadweight loss, which is highly sensitive to δ, we present deadweight loss as a percentage of the buyer's value, which is stable across different choices of δ.

5.2. **Estimates.** Table 2 displays point estimates and standard errors for the market tightness parameters, the per-period bidding cost, and the per-period measure of new entering agents. Note that λ_2 is significantly greater than 0 at the 1% level, which means our estimates reject the assumption that auction size follows a standard Poisson distribution. In particular, the Poisson model substantially underestimates the dispersion in the number of bidders matched to a given auction.

Figure 4 depicts point estimates for our Stage 1 distribution estimators (thick, solid lines), point-wise confidence bounds for a selected grid of domain points (vertical box plots), and the 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 market sample selection process would lead to a very different view of the distribution of auction participation. This substantial difference shows up in the mean of $\tilde{K}$ and K (4.07 versus 7.96 respectively) and the variance (3.19 versus 14.46 respectively). 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 $H\left(y; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b\right)$ 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 4. Stage 1 Estimates

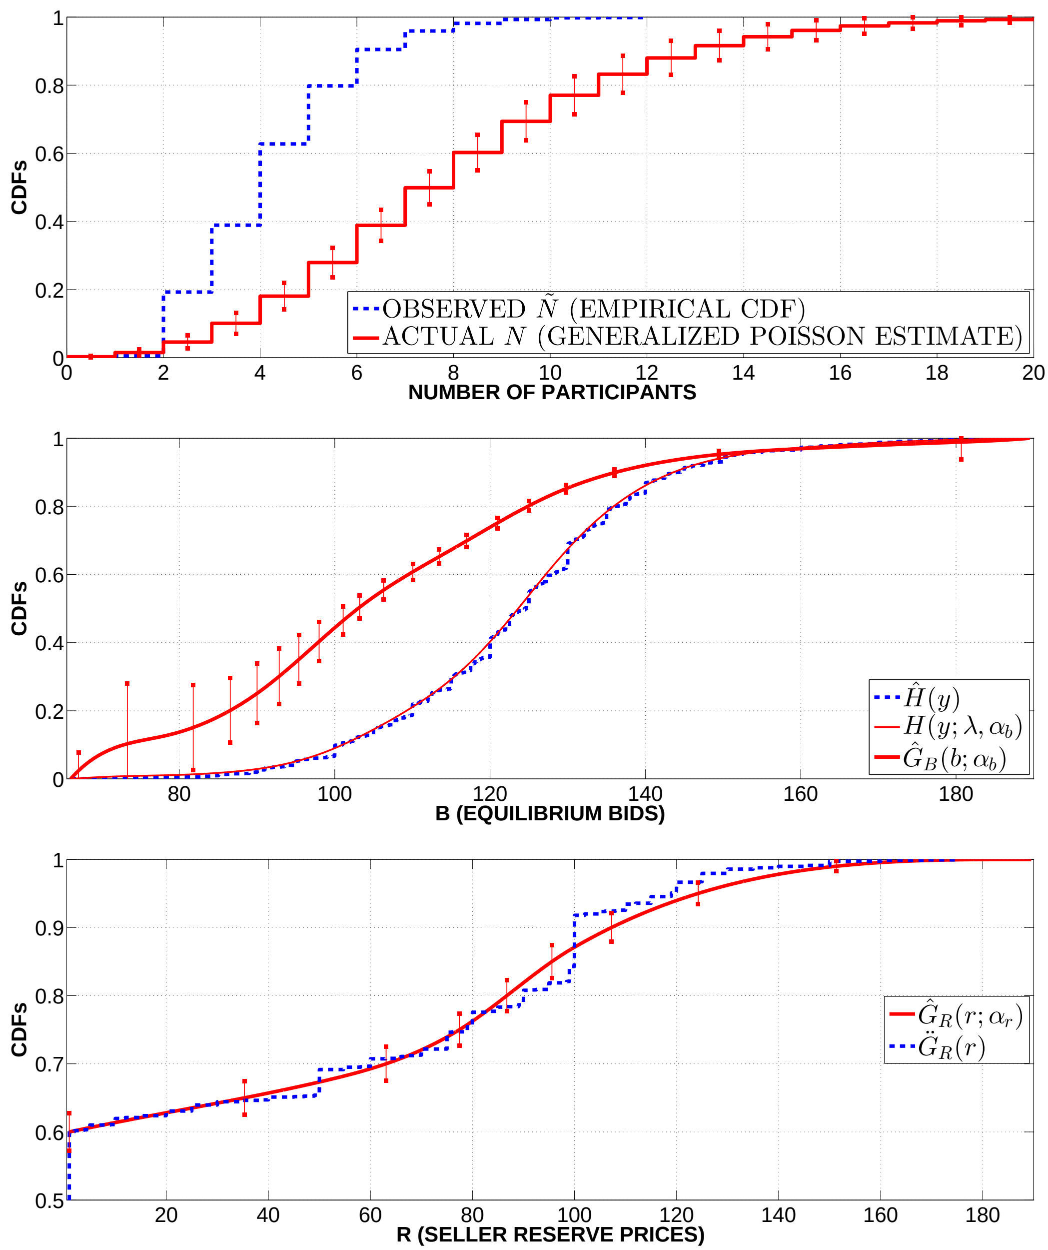

Figure 5 presents the dynamic inverse bid functions $\hat{\beta}^{-1}\left(b; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b, \hat{\boldsymbol{\alpha}}_r, \delta\right)$ which we estimate for a grid of values of the time discount factor δ between 0.75 and 0.98. Recall from Figure 1 that the 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 $\boldsymbol{\lambda}$, 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

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

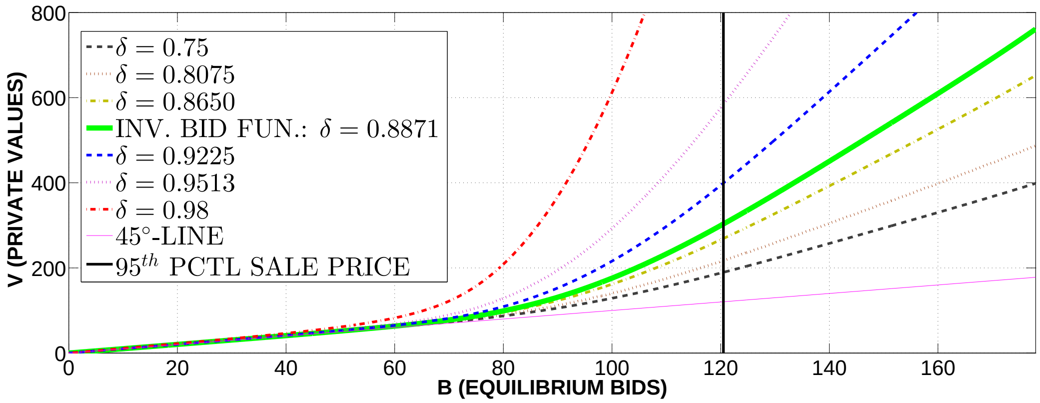

TABLE 3. Mean Private Values and Information Rents For Various δ

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%

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 95^{th} percentile of the private value distribution is over \$1,300, which we consider to be implausibly high. When required to show a counterfactual result pertinent to a particular value of δ (*e.g.*, Figure 1)[33], we will use $\delta = 0.8871$ as it is the median of the set of δ values we consider. However, as we move on to welfare calculations in our counterfactual analysis, we will focus discussion predominantly on measures that are stable across choice of δ.

Table 3 displays various descriptive statistics derived from the Stage 2 estimates, including average private values, average private values of winners, and information rents defined as the difference between the winner's private value and the spot-market price.

[33]The fact that dynamic demand shading incentives are more important than the static demand shading incentives is not sensitive to the choice of δ within the range we consider.

FIGURE 6. Type Density Estimates for $\delta = 0.8871$

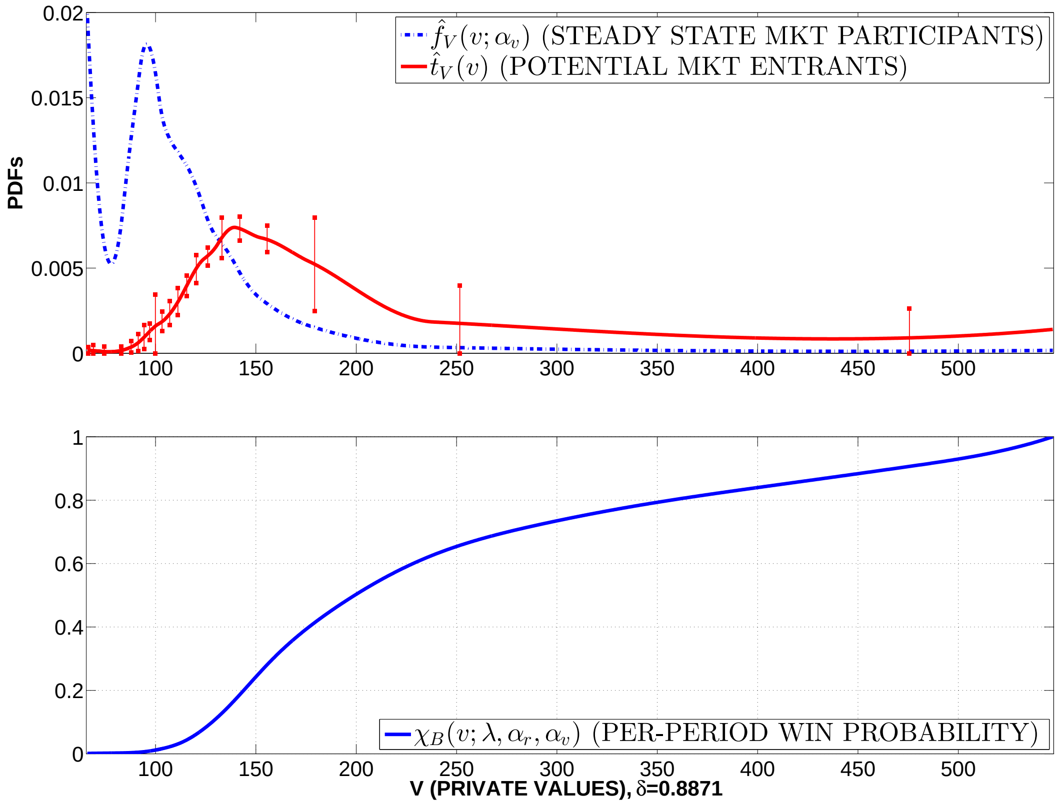

The last row of the table shows information rents as a fraction of the winner's private value, on average. Finally, Figure 6 presents other Stage 2 estimates related to the distribution of buyer values. The upper pane displays the PDF of the distribution of market participants' private values in steady state under our preferred specification, $\hat{f}_V(v; \boldsymbol{\alpha}_v, \delta = 0.8871)$ (dash-dot line) and the type distribution for new market entrants each period, $\hat{t}_V\left(v; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b; \hat{\boldsymbol{\alpha}}_r, \delta = 0.8871\right)$ (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 buyer value. 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. On a related note, the delay between entry and trade implies that the lower-value buyers are the ones most affected by the per-period participation cost, $\hat{\kappa} = \$0.065$.

In comparing estimates for f_V and t_V in Figure 6, two important differences should be noted. First, f_V depicts the steady-state type distribution of the measure $\lambda_1/(1-\lambda_2) = 7.96$ of potential entrants in each period. On the other hand, t_V represents the type distribution of the measure $\mu = 0.9649$ of new potential entrants added to the market each period in order to maintain the steady state. 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 a success. On the other hand, t_V depicts a selected set of buyers who move in and out of the market each period. t_V has a higher density of high value agents than f_V as these high value agents are more likely to transact and exit the market each period than are low value agents.

6. COUNTERFACTUALS

We now perform a series of three counterfactual analyses to investigate the economic implications of our structural model. The first explores market efficiency concerns. The second decomposes the relative importance of what we refer to as *platform composition* (PC) effects (*i.e.*, market entry/exit when general 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 reserve price design in the dynamic market setting where a seller may also return in a future period to re-list his good if it does not sell today. For notational simplicity we omit the parameter arguments of structural primitives unless needed for clarity. Appendix C in the online supplement contains an overview of the algorithm used to compute equilibria for the structural counterfactuals presented in this section.

6.1. **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. Recall that even when the spot-market mechanism is efficient within a given auction, inefficiency may arise if multiple high value buyers are matched to a particular auction. Our welfare calculations are meant to capture the frequency and magnitude of these inefficiencies.

6.1.1. *"Model Anemic" Inefficiency Calculations.*

In this section we use only our Stage 1 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. Rather, 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 between high-value buyers that ought to receive the good in an efficient allocation and lower-value buyers that ought not. Since the buyer-seller ratio is $\lambda_1/(1-\lambda_2)$, the efficient cutoff in private value space is defined by $v_{eff} \equiv F_V^{-1}\left(1 - \frac{1-\lambda_2}{\lambda_1}\right)$. We can re-define this cutoff in bid space (where the raw data live) as $b_{eff} \equiv G_B^{-1}\left(1 - \frac{1-\lambda_2}{\lambda_1}\right)$.

Since we only observe the highest losing bid within each auction, we can only provide a lower bound on the frequency of inefficient allocations using raw data and our cutoff. Intuitively, if the highest losing bid in a given auction exceeds b_{eff}, then the corresponding losing bidder would have received the good under an efficient allocation. We find that 28.47% of the auctions in our sample end with at least one losing bid above the efficient cutoff.[34] 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}. [35] 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 the private values that underpin observed bids.

6.1.2. *Structural Welfare Calculations.* Our full Stage 2 structural estimates allow us to get a more complete idea of the frequency and magnitude of market inefficiency. Recalling that the probability that an entrant of type v wins an auction is $\chi\left[\beta(v)\right]$ and that the mass of this type of buyer[36] is $f_V(v)\mathcal{C}^\infty/\mathcal{S}^\infty$, then we can write the probability of an inefficient allocation as

$$\Pr[V_{winner} < v_{eff}] = \int_{\underline{v}}^{v_{eff}} \chi\left[\beta(v)\right] f_V(v) \frac{\mathcal{C}^\infty}{\mathcal{S}^\infty} dv.$$

Note that this measure is invariant to the choice of the time discount factor δ. Our point estimates imply that 35.89% of Kindle auctions on eBay end with an inefficient outcome. Deadweight loss calculations in levels will be sensitive to choice of δ. In order to address this problem, we adopt a measure that we refer to as the *efficiency ratio*

$$\mathcal{E}_{u,\delta} = \frac{\int_{\underline{v}}^{\overline{v}} s\chi_u(\beta_u(s)) f_{V,u}(s) ds}{\int_{v_{eff}}^{\overline{v}} s f_{V,u}(s) ds}$$

where 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

[34] There 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.

[35] 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.

[36] Recall that the measure of sellers is normalized to 1, meaning the total measure of potential entrants is the buyer-seller ratio $\mathcal{C}^\infty/\mathcal{S}^\infty$.

allocation. The u subscript denotes number of units involved in each auction listing for our counterfactual centralization analysis below; for now, we simply fix $u = 1$. By expressing surplus as a fraction of total possible surplus, the influence of δ in the numerator and denominator largely cancels out and we get a measure that is much more stable across different assumptions on time discounting (see alternative calculations displayed in the first row of Table 4). Although the efficiency ratio is (by construction) between zero and one, we also compute the efficiency ratio under a hypothetical lottery system, denoted $\mathcal{E}_{lott,\delta}$, as the relevant minimum efficiency benchmark (see last row of Table 4).

With these definitions in hand, our point estimates imply that the fraction of total deadweight loss is simply $1 - \mathcal{E}_{1,0.8871} = 0.135$ under our preferred specification. To put this number into context, deadweight loss under a lottery system (see Table 4) 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 (or a random lottery) 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.

6.1.3. *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 homogeneous products and all transactions are covered by eBay's consumer protection program, we think it is reasonable to assume that buyers view them as nearly perfect substitutes for one another. This leads us to think 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 setting with $u \geq 2$. First, each u-unit auction attracts a number of bidders K_u distributed as a generalized Poisson random variable with expected value

$$\mathrm{E}[K_u] = \frac{\lambda_{1,u}}{1-\lambda_2} = u\frac{\mathcal{C}_u^\infty}{\mathcal{S}^\infty}. \tag{31}$$

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**	**Buyer-Seller Ratio** $\mathcal{C}_u^\infty / \mathcal{S}^\infty, \delta = 0.88$
1	0.89	0.88	0.87	0.86	0.85	0.84	0.82	7.75
2	0.92	0.92	0.91	0.91	0.91	0.90	0.89	5.50
4	0.94	0.94	0.94	0.94	0.94	0.93	0.93	3.95
8	0.95	0.95	0.95	0.95	0.95	0.95	0.95	2.92
Lottery	0.58	0.54	0.49	0.47	0.41	0.35	0.26	7.75 (assumed)

which is just the ratio of buyers to u-unit auctions. We assume that λ_2, the dispersion parameter, is fixed at the estimated value and allow $\lambda_{1,u}$, the size parameter, and $\mathcal{C}_u^\infty$, the measure of buyers, to adjust so that (31) 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 that applies for all u units being auctioned. The use of a reserve price makes it possible to make comparisons between the status quo setting and the uniform price setting we study here. Each bidder submits a bid to the auction at which he or 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.

Throughout this exercise we hold fixed the structural primitives such as the per-period bidding cost κ, the measure of new potential entrants entering the game each period μ, and the distribution of new potential entrant types t_V. When we compute our counterfactual equilibria, we let $\underline{v}_u$, $\lambda_{1,u}$, $F_{V,u}$, and $f_{V,u}$ adjust to satisfy our equilibrium conditions.[37]

One of the general takeaways from this research is that understanding the impact of platform market design on the participation decisions of agents 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 limit model allows us to handle these questions by computing the counterfactual, steady-state SCE given a platform based on a u-unit spot-market mechanism. Table 4 provides results for counterfactual efficiency ratio statistics for $u \in \{1, 2, 4, 8\}$. Recall from above that the efficiency ratio

[37] In all of the counterfactuals we present, $\underline{v}_u$ is at least as large as in the data-generating process.

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

compares gains from trade in a single period of a u-unit auction platform market with the welfare generated by an efficient allocation of the goods within that period.

We would like to draw attention to two features of our results. First, welfare ratios are remarkably stable within each alternative market structure across different specifications of the time discount factor δ. Second, although the eBay platform fails to realize a significant portion of possible gains from trade, it is very close to an alternative platform structure that reaps nearly maximal social benefit. The majority of possible gains from centralization can be realized by only 2- or 4-unit uniform-price auctions, so there is little need shift toward a fully centralized market.

We have included the buyer-seller ratio to get a sense of the impact of centralization on the composition of the bidders.[38] There are two effects of note. First, as the market centralizes, the low value bidders exit the market. Second, as the market centralizes the ratio of buyers to sellers drops. The net effect is that as the market centralizes, the buyers face fewer competitors, but the competitors are significantly stronger.

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. However, a careful examination of the moving parts within the model indicates that the sign of this effect is ambiguous. On the one hand, PC effects arise from buyer exit at the low end of the type distribution. On the other hand, complicated DI effects ensue: a given bidder above the new participation cutoff $\underline{v}_u$ faces fewer overall competitors in the market, but her remaining competitors have higher values on average. These complex interactions make it difficult to derive predictions on bidding behavior and the resulting effects of revenue. Table 5, which contains the mean revenues generated per-auction as a function of u, demonstrates a somewhat counterintuitive result that the average sale price actually *falls* as u increases.

[38] We have included the buyer-seller ratio for the $\delta = 0.8871$ case, but the ratios for other choices of delta differ by less than 3%.

FIGURE 7. The Efficiency-Revenue Link

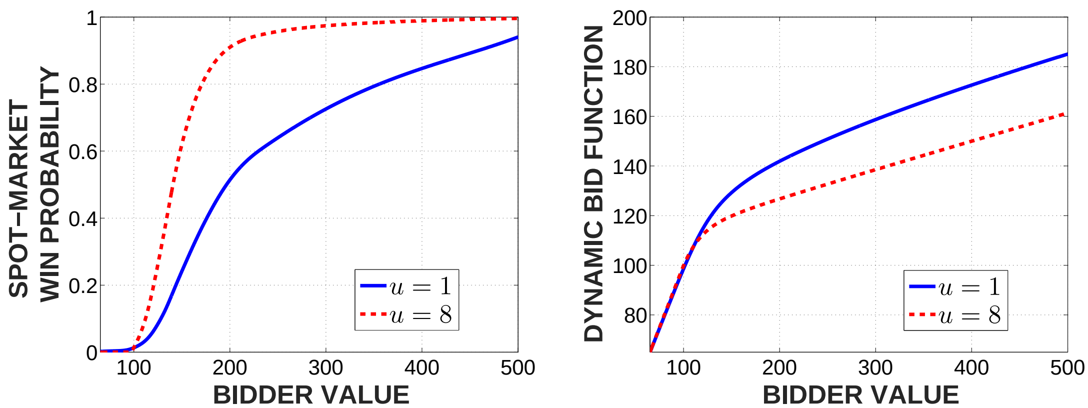

To help explain why revenue drops as efficiency rises, Figure 7 plots the probability of winning for each type of agent and the equilibrium bid function for the $u = 1$ (solid line) and $u = 8$ (dashed line) market structures.[39] The win probability plot reveals that for most agents (especially those 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, which 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 charges commissions 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 aggregate 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 $\kappa/\chi_u(\beta(v))$, which ought to be interpreted as the average participation cost paid by an agent of type v before winning an auction. The ALPC refers to this quantity averaged over the steady-state distribution of buyers multiplied by the buyer to seller ratio:

$$\mathcal{ALPC}_{u,\delta} = \frac{\mathcal{C}_u^\infty}{\mathcal{S}^\infty} \int_0^\infty \frac{\kappa}{\chi_u(\beta_u(v))} f_{V,u}(v)dv$$

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] each participant pays on average a total of \$8.98 each over his or her lifetime in the market. As one would expect, these costs are disproportionately

[39]The probability of winning in the $u = 8$ market is truncated, which reflects the fact that the infra-marginal market participant bidder has a higher value under the $u = 8$ market structure.

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

TABLE 6. Aggregate Lifetime Participation Costs $\mathcal{ALPC}_{u,\delta}$

# 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

borne by the lower value agents that must enter the market repeatedly to win an item. Because many of these low value agents build up in the steady-state type distribution, the aggregate lifetime participation cost is very high relative to κ. Expected lifetime participation costs for winners in a period are much smaller as these are more likely to be high value entrants that exit the market quickly. When $\delta = 0.88$ and $u = 8$, each participant pays on average a total of \$9.50 over the course of his or her participation in 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.

6.2. **Relative Importance of Platform Composition and Dynamic Incentives.** Our goal in this section is to measure the relative importance of the DI and PC effects. As an illustrative example, we consider increases to the per-period participation cost κ. Aside from illuminating answers to questions of academic interest, this counterfactual may also provide practical guidance to eBay and other online market designers regarding which 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 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 PC effects where low-value buyers exit the market, which reduces the buyer-seller ratio and strengthens the steady-state distribution of active bidder types.

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 decompose 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 value function (which drives the DI effect) or the buyer-seller ratio and bidder value distribution (which drives 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. We denote endogenous objects in an equilibrium with cost κ using the subscript κ.

Let $\mathcal{R}_\kappa$ denote the ratio of entrants to sellers in an equilibrium with participation cost κ.[41] Note that $f_{V\kappa}$ and $F_{V\kappa}$ live on support $[\underline{v}_\kappa, \overline{v}]$ with $\underline{v} \leq \underline{v}_\kappa$ whenever $\underline{\kappa} < \kappa$. Since we will be conducting out-of-equilibrium comparisons, we re-define the probability of a buyer winning as

$$\chi_\kappa(v; \beta_\kappa, \boldsymbol{\lambda}_\kappa, F_{V\kappa}) = G_R(\beta_\kappa(v)) \sum_{m=0}^{\infty} \pi_M(m, \boldsymbol{\lambda}_\kappa) F_{V\kappa}(v)^m$$

The first term 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 the probability a bidder with value v wins an auction in a given period in a steady-state SCE with participation cost κ.

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

$$\mathcal{W}(\beta_\kappa, \boldsymbol{\lambda}_\kappa, F_{V\kappa}) \equiv \mathcal{R}_\kappa \int_{\underline{v}}^{\overline{v}} s \chi_\kappa(s; \beta_\kappa, \boldsymbol{\lambda}_\kappa, F_{V\kappa}) f_{V\kappa}(s) ds,$$

where for convenience we 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 allocative efficiency is the *dynamic gap*

$$\mathcal{DG}(\underline{\kappa}, \kappa) \equiv \mathcal{W}(\beta_\kappa, \boldsymbol{\lambda}_{\underline{\kappa}}, F_{V\underline{\kappa}}) - \mathcal{W}(\beta_{\underline{\kappa}}, \boldsymbol{\lambda}_{\underline{\kappa}}, F_{V\underline{\kappa}}).$$

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 bidding function from an equilibrium generated by a participation cost $\kappa > \underline{\kappa}$. The idea is to hold fixed the endogenous quantities that correspond to PC effects ($\boldsymbol{\lambda}$ and F_V) while allowing DI effects (the bidding function) to vary with κ.

The *platform gap*, $\mathcal{PG}$, captures the role of PC effects in allocative efficiency

$$\mathcal{PG}(\underline{\kappa}, \kappa) = \mathcal{W}(\beta_{\underline{\kappa}}, \boldsymbol{\lambda}_\kappa, F_{V\kappa}) - \mathcal{W}(\beta_{\underline{\kappa}}, \boldsymbol{\lambda}_{\underline{\kappa}}, F_{V\underline{\kappa}})$$

This gap is computed by comparing equilibrium allocative efficiency generated by $\underline{\kappa}$ to an out-of-equilibrium market with the same bidding function but matching parameters and steady-state distributions of an equilibrium with a higher cost κ. Here we hold DI

[41] Including notation for both $\mathcal{R}_\kappa$ and $\boldsymbol{\lambda}_\kappa$ is not strictly necessary since the former can be computed from the latter. However, we do so here for expositional clarity.

FIGURE 8. Relative Size of Dynamic and Selection Gaps

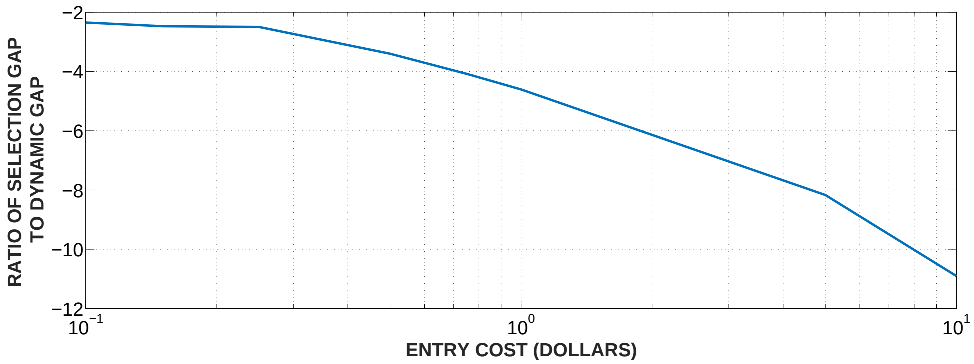

effects (the bidding function) fixed and vary endogenous quantities that correspond to the PC effect ($\boldsymbol{\lambda}$ and F_V).

As a first comment, the allocative efficiency must rise as bids rise since the bids are more likely to surpass reserve prices. As higher values of κ are associated with lower continuation values, and hence higher bids, the DI effect always increases allocative efficiency as κ rises (*i.e.*, $\mathcal{DG}(\underline{\kappa}, \kappa) > 0$ if $\kappa > \underline{\kappa}$). Although the PC effect is of ambiguous sign in theory, the PC effect is negative in our data. 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 powerful 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 platform composition effects of market changes can be many times more important than understanding the dynamic incentive effects of the changes.

6.3. **Optimal Reserve Prices.** As has been regularly noted about the eBay market place, sellers tend to choose low reserve prices. In our data, almost 60% of the reserve 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 reserve price equal to $\beta(\underline{v})$, the lowest possible equilibrium bid in the auction.[42]

Our goal is to assess the strength of the incentives of the sellers to carefully choose a revenue maximizing reserve price. To place a number on this we consider the benefits to a single seller of optimally choosing his or her reservation price relative to choosing a reserve price of \$0.99. 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 $\boldsymbol{\lambda}$, F_V, and (e, β) at their

[42]Such a reserve 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.

status quo values. The problem the seller solves is:

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

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 reserve price varies from a low of \$84.90 to a high of \$85.80. At the optimal reserve price, the revenue generated is either \$122.30 or \$122.31 across all of the possible δ. This represents an improvement of just \$0.95 relative to a reserve price of \$0.

The benefits from optimally choosing the reserve price are small because each seller is matched with 7.96 bidders in expectation, which means that the competition between bidders will be intense. Bulow and Klemperer [1996] show that choosing the reserve price optimally is no better than adding a single extra bidder to the market. With almost 8 bidders on average already participating, it should not be surprising that there is little room left for an optimal reserve price to have a significant effect on auction revenues.

7. CONCLUSION

Our goal has been to provide a model of a marketplace that is both rich enough to capture the salient features of the market (*e.g.*, the large number of auctions concluding each day and the potential costs of participation) and yet remain tractable enough to facilitate empirical analysis using commonly available observables from platform markets. 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. Moreover, we have also demonstrated that the structural components of this simplified 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 a variety of pricing rules. Finally, we have also proposed a simple but flexible GMM estimator to allow researchers to estimate the structural primitives and, in turn, compute counterfactuals.

One immediate conclusion of our analysis relates to the importance of intertemporal incentives. In online auction markets bid shading driven by the opportunity cost of winning today (in light of further buying opportunities tomorrow) is likely to be much larger in magnitude than the more commonly studied static bid shading incentives due to non-truthful pricing mechanisms. This dynamic demand shading 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.

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, their sheer size may also give rise to search frictions which prevent market outcomes from attaining the social ideal. Our model estimates within the context of the market for Kindle Fire tablets facilitates an exploration of residual inefficiency in eBay's decentralized, single-unit auction platform. 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 by running multi-unit, uniform-price auctions with as few as 4 units each, 2/3 of the welfare loss can be recovered.

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 value 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.

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. 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.

Jorge Balat. Highway procurement and the stimulus package: Identification and estimation of dynamic auctions with unobserved heterogeneity. *Typescript, Johns Hopkins University*, 2013.

Patrick Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, Inc, 1968.

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.

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.

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. *typescript, 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.

Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71:1443–1489, 2003.

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. Optimal auction design. *Mathematics of Operations Research*, 6:55–73, 1981.

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.

Brennan C. Platt. Inferring ascending auction participation from observed bidders. *Typescript, 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.

Maher Said. Sequential auctions with randomly arriving buyers. *Games and Economic Behavior*, 73 (1):236–243, 2011.

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

Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, XLIII:462–476, 2006.