HOW EFFICIENT ARE DECENTRALIZED AUCTION PLATFORMS?

AARON L. BODOH-CREED - UNIVERSITY OF CALIFORNIA, BERKELEY
JÖRN BOEHNKE - HARVARD UNIVERSITY
BRENT HICKMAN - UNIVERSITY OF CHICAGO

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

Keywords: Dynamic Auctions, Approximate Equilibrium, Internet Markets.

JEL Codes: C73, D4, L1

Date: October 2017.

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

1. INTRODUCTION

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

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

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

On the empirical front, we make several contributions to the literature on identifying auction models. A key feature of our estimator is that it requires only observables that are readily available on many platform websites. In particular, we are able to identify the average number of buyers per auction without assuming that we observe all of the bidders in each auction. If bid submission times are randomly ordered, then some auction participants with an intent to bid may be prematurely priced out of the spot-market before they get a chance to submit their bid. Therefore, the total number of unique bidders within

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

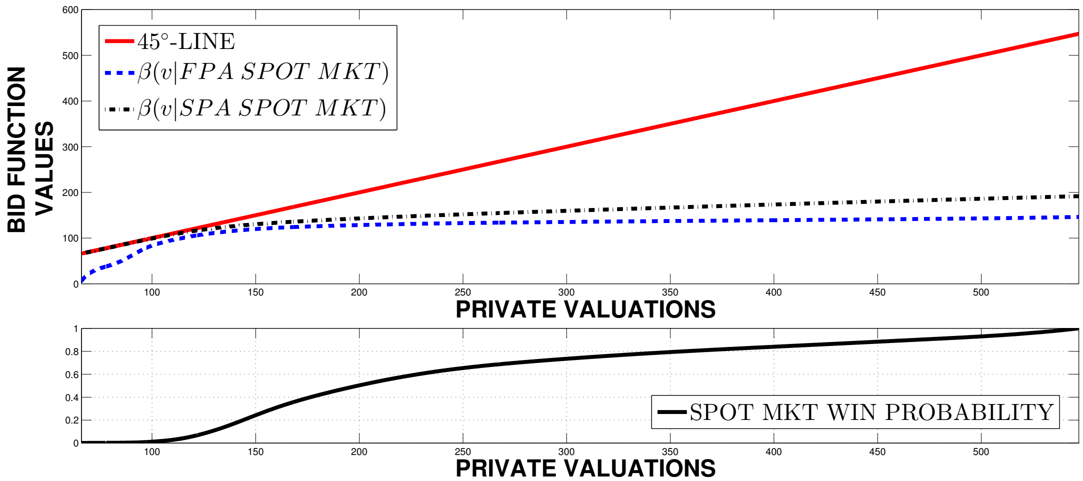

FIGURE 1. Static Versus Dynamic Demand Shading Incentives

a given eBay auction constitutes a lower bound on the actual number of competitors. Our nonparametric identification argument for the dynamic structural model requires only that we observe this lower bound on the number of competitors, the seller reserve price, and the highest losing bid within each auction.

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

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

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

Having estimated the structural model, we move on to investigate the efficiency of the market. Platform markets like eBay exist for the purpose of reducing frictions that impede trade and converting some of the resulting efficiency gains into profits for the platform's owners. With this in mind, it is natural to assess how closely eBay approaches the ideal of fully efficient trade. To fix ideas, suppose there are exactly two listings and four bidders, two bidders with high valuations and two bidders with low valuations. The social planner's preferred outcome is one where each auction attracts one high value bidder as this would guarantee that the high value bidders win in any monotone bidding equilibrium. However, when there is randomness in the bidder-listing match process, there will be a positive probability that one auction listing will not have a high-value participant, meaning a low value bidder wins at a low price and a high-value bidder loses. Another way of putting it is that the matching frictions mean some auctions have too much competition and others too little relative to an efficient allocation.

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

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

Next, we explore the implications of alternative spot-market mechanisms eBay could use to improve efficiency. Specifically, we consider the welfare cost of eBay's choice to use single-unit auctions, which we refer to as *decentralization*. We use our estimates to analyze outcomes of alternative markets where, instead of single-unit auctions, eBay runs K-unit, uniform-price auctions, and we use K as a measure of the market's *centralization*. The most extreme version of this counterfactual would be a single multi-unit, uniform-price auction each day. We find that aggregating auctions together so that eBay runs half as many auction listings for 2-units each day recovers 35% of the welfare loss by improving the efficiency of the allocation, while running a quarter as many 4-unit auctions recovers over 57% of the welfare loss. However, centralization reduces the expected sale price, which in turn reduces eBay's revenues. In addition to being a vehicle for analyzing the welfare losses, we believe that centralizing auctions is a practical design strategy in settings wherein the goods are homogenous (e.g., new Kindles).

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

Our last counterfactual considers the relative impact of dynamic incentives (i.e., bid shading driven by opportunity costs) and the composition of buyers using the platform. We analyze the relative impact of these two forces in a counterfactual study of the effects of an increase in the participation cost for bidders. An increased participation cost pushes low-value bidders out of the market (i.e., a platform composition effect), altering the steady-state distribution of bidder values and the ratio of buyers to sellers. The cost increase also has an effect on dynamic incentives by reducing the continuation values (and hence raising the bids) of the buyers. We find that the platform composition effects have

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

from two to ten times more effect on market efficiency than the dynamic incentive effects. Although we consider a particular change in the market structure, our analysis emphasizes the importance of attending to the selection of users into the platform when redesigning any platform market.

Finally, our paper also makes a contribution to the theory underlying the large market models we use. The notion of a large market approximation, sometimes referred to as an Oblivious Equilibrium, is not novel to this paper. However, proving a formal relationship between a model with a continuum of players and the finite markets that exist in reality is difficult when the market mechanism admits discontinuities, and an auction setting provides several points where such discontinuities can arise. In Appendix B, we prove that despite these issues, one can view an equilibrium of the model with a continuum of buyers and sellers as an ε-equilibrium of the model with a finite number of buyers and sellers. We view this result as a justification for our use of the continuum model in our estimation and counterfactual exercises.

The remainder of this paper has the following structure. In Section 2 we develop a theory of bidding within a dynamic platform based on a model with a continuum of buyers and sellers. In Section 3 we use this model to specify a parsimonious structural model of eBay, which we show is identified from observables. We also propose a semi-nonparametric estimator based on B-splines. In Section 4 we present our model estimates. Section 5 presents our counterfactuals on welfare and the relative importance of platform composition and dynamic incentive effects. Proofs for the technical claims in the main text are relegated to Appendix A. In Appendix B, we prove that our model with a continuum of agents approximates an analogous model with a large, but finite, number of participants. Appendix C provides details on our estimation techniques and algorithms. Appendix D provides a variety of additional counterfactuals on revenue, participation costs, and seller incentives.

1.1. **Related Literature.** The most closely related paper to ours is the contemporaneous Backus and Lewis [2016], which studies a model of eBay where bidders participate in a sequence of single-unit, second-price auctions. Backus and Lewis [2016] focuses on identifying a buyer demand with flexible substitution patterns across different goods and the possibility that individual bidder demand evolves over time. Their model could be used to compute welfare counterfactuals like ours in a setting with homogenous goods, but this (obviously) excludes substitution between products. However, since Backus and Lewis [2016] does not include participation costs or a procedure for estimating the measure and type distribution of buyers entering the market, the model remains silent on how market structure influences entry and exit through the interaction of changing auction format and participation costs. Our structure also decomposes static and dynamic demand shading incentives, which allows us to compare the relative strengths of these forces. This decomposition also links existing results on estimating static auction models with the literature on dynamic auction markets. Finally, we prove that our large market model approximates a more realistic model with a finite number of agents. Due to the different foci and contributions of each work, we view our papers as complementary.

Our methodology analyzes approximate equilibria played by a large number of agents, which has been a prominent theme in the microeconomics and industrial organization literatures. Due to the broad scope of this literature, we provide only a brief survey and a sample of the important papers related to the topic. Early papers focused on conditions under which underlying game-theoretic models could be used as strategic microfoundations for general equilibrium models (e.g., Hildenbrand [1974], Roberts and Postlewaite [1976], Otani and Sicilian [1990], Jackson and Manelli [1997]). Other early papers focused on conditions under which a generic game played by a finite number of agents approaches some limit game played by a continuum of agents (e.g., Green [1980], Green [1984], and Sabourian [1990]). A more recent branch of this literature applies these ideas to simplify the analysis of large markets with an eye to real-world applications (e.g., Fudenberg, Levine, and Pesendorfer [1998]; MacLean and Postlewaite [2002]; Budish [2008]; Kojima and Pathak [2009]; Weintraub, Benkard, and Roy [2008]; Krishnamurthy, Johari, and Sundararajan [2014]; and Azevedo and Leshno [2016]).

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

The estimation of dynamic auctions is a relatively new field that is maturing rapidly. The first paper we are aware of that attempts to estimate a model of a sequence of (procurement) auctions is Jofre-Bonet and Pesendorfer [2003]. This paper estimated the effect of capacity constraints on bidder behavior and the efficiency of the auction outcomes. In a later paper, Donald, Paarsch, and Robert [2006] estimated a model of a sequence of timber auctions wherein bidders have multi-unit demand, and the focus of this second paper was price patterns across auctions. Neither paper solved for counterfactuals due to the computationally intractable models used in the estimation exercises.

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

2. A MODEL OF PLATFORM MARKETS

Before describing our formal model of bidder behavior, we would like to informally describe the behavior our model is intended to capture. Our informal description is not meant to be universal, but we believe it to be fairly typical of behavior on eBay. We imagine that before entering the eBay market, the buyer considers her own value for the good, makes an assessment of the opportunity cost of winning, and formulates her bid. After entering the market, the bidder considers bidding in an auction that is closing in the near future. We assume that the time a bidder chooses to enter the market is driven by factors exogenous to eBay (e.g., the schedule of work breaks), which means the buyer only considers bidding on a small and randomly selected fraction of the auctions that close during that day. We refer to the mechanism used in individual auctions as the *spot-market mechanism*. If a buyer wins the spot-market auction, then she does not participate in future days. Remaining buyers return to eBay the next day to place a bid. We summarize the timing in the following diagram.

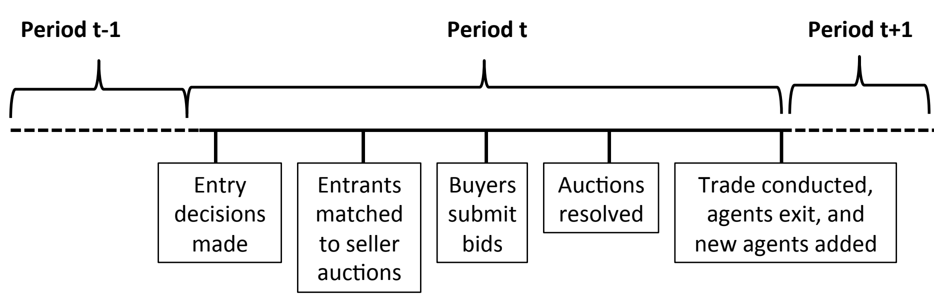

FIGURE 2. Timing within a period

Let us now address a few potential objections to our informal story. First, is it reasonable to assume that the buyer formulates her bid before entering the market? This behavior would be fully rational if eBay used a second-price auction (SPA) format. However, since eBay auctions use a non-SPA pricing rule, there are static bid shading incentives that push the bidder to adjust her bid based on the competitive environment (e.g., the current highest bid) she observes upon entering an auction. However, as discussed in the introduction (see

Figure 1) and in Section 4, intertemporal incentives that are independent of the auction she participates in today play a far more important role in determining the optimal bid than the competitive environment of a particular auction. Therefore the bidder has relatively weak incentives to update her bid according to the details or history of any given auction.

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

Third, our informal story assumes the bidder does not repeatedly bid within a particular auction. As we discuss more thoroughly in Section 3, we only consider bids arriving near the end of the auction. One reason we do this is to avoid the question of how to handle dynamic behavior within a given auction. As it stands, we see relatively few instances in the data of a bidder returning to increase her bid if she has already bid within the narrow window we consider at the close of each auction.

Now we formalize the scenario we described above. Whenever possible, we develop theoretical results in terms that apply for arbitrary, well-behaved spot-market pricing rules so that our model may serve as a general framework for quantitative analysis of platform markets. Our definition of "well-behaved" is captured by Assumption 2.5 below.

2.1. **Model Primitives.** We treat sellers as a source of exogenous supply, and we assume their decisions (e.g., entry/exit and starting prices) are exogenous. This modeling choice is driven by the fact that sellers face very weak incentives to set the optimal starting price,[3] which makes estimates based on a model of seller behavior less credible. Therefore, we present a theory of the buyer side of the market.

Although our model is dynamic, our analysis focuses on stationary equilibria where the aggregate states that characterize the market are constant across periods in equilibrium. This also means that the strategies describing an agent's actions explicitly condition on the agent's private information and implicitly condition on the equilibrium values of the aggregate state variables. We discuss issues such as how to extend our model to the nonstationary case and what is lost by assuming stationarity in Appendix B, in which we prove that our stationary equilibrium is an approximate equilibrium of a model with a finite number of agents.

[3]In Appendix D, we show that the sellers earn less than $1 in increased revenue by moving from a starting price of $0 to the revenue maximizing starting price.

The market evolves in discrete time with periods indexed $t \in \{0, 1, 2, ...\}$. In each period there is a measure 1 of sellers with starting prices described by cumulative density function (CDF) G_R with support $[0, \bar{r}]$.[4,5] G_R may have a mass point, but only at the lowest possible reserve price, $r = 0$, and has a probability density function (PDF) $g_R(R|R > 0)$ that is strictly bounded away from zero over the rest of its support.

We refer to the set of buyers present at the start of period t as *potential entrants*; at the beginning of each period they make decisions based on the observed number and type distribution of the other potential entrants and their own types. Each period, the first choice each potential entrant must make is whether or not to enter the market and participate in the platform. We denote the choice to participate as *Enter* and refer to the agents that make this choice as *entrants*. The choice to not participate is denoted *Out*, and if the agent chooses this option she leaves the game permanently.

Throughout we assume that the goods for sale are homogenous and that buyers have demand for a single unit.[6] Each buyer's value for the good is her private information, which we denote as v and assume is fixed over time. A buyer that wins a good on the eBay platform and pays a price of p receives a payoff in that period of $v - p - \kappa$, where κ is a per-period participation cost paid by entrants regardless of whether they win. We assume $\kappa > 0$; this may reflect the opportunity cost of time spent searching for a listing and participating in the market, or it may reflect an actual monetary participation fee that the platform designer imposes. 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.

Each period there is a measure C continuum of potential entrants with a type distribution equal to F_V.[7] A measure of potential entrants equal to μ is added to the economy at the end of each period, and the distribution of the values of these new potential entrants has CDF $T_V(\cdot)$ with PDF $t_V(\cdot)$. We assume that t_V is strictly positive over the support $[0, 1]$. F_V includes both newly added potential entrants and ones remaining from the previous period and is an element of the space of probability measures over $[0, 1]$, denoted $\Delta([0, 1])$. Unless stated otherwise, $\Delta([0, 1])$ is endowed with the Kolmogorov topology, which is metrized by the sup-norm over the space of CDFs associated with the measures in $\Delta([0, 1])$ (see Appendix A). A generic, measure 0 buyer is denoted using the subscript i. The state variables describing the economy in a given period is a vector (C, F_V, G_R).

[4]The lowest starting price on eBay is \$0.99, but this does not affect our theoretical results. eBay also allows sellers to choose reservation prices that are hidden from buyers, but this is done so infrequently that we ignore it in our modeling.

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

[6]We discuss the homogeneity of the goods in our data set in Section 4.

[7]Since the letter β is used later on to denote bids, we chose C, for "consumer."

After choosing *Enter*, each entrant formulates a strategic bid in a a simultaneous-move spot-market without knowing either the number or identity of the other agents participating in the particular auction to which she is matched. The form of the random matching process, the distribution of entrant types, and the exogenous distribution of starting prices is known to agents at the point when they choose their bids. If a measure $\mathcal{C}$ of buyers chooses to enter the auction market (out of a measure C of potential entrants), the entrants are randomly assigned to auctions with each auction receiving a random number K of bidders where $Pr\{K = k\} = \pi(k, \boldsymbol{\lambda})$. 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 refer to $\boldsymbol{\lambda}$ as the *market tightness* parameter since it is determined by the buyer-seller ratio on the platform, $\mathcal{C}$, in our application. The key properties we assume for the bidder arrival process are:

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

(1) $E[K] = \mathcal{C}$
(2) $\pi(k = 1, \boldsymbol{\lambda}) > 0$ *for all* $\boldsymbol{\lambda}$
(3) *π is continuous in $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}$ is continuous in $\mathcal{C}$ where continuity is with respect to the product topology on* $\mathbb{R}^{\infty}$

Part (1) requires that all of the buyers be matched into auctions. If we allowed $E[K] > \mathcal{C}$, then in aggregate the auctions would be assigned more bidders than entrant buyers that exist in the market. If we allowed $E[K] < \mathcal{C}$, then a positive measure of bidders would not be assigned to any auction. If the market truly is large, then an auction must be closing in the near future, and one wonders why each unassigned buyer chose not to participate in one such auction. Part (2) insures that an entrant with the lowest value amongst all entrants can win an auction with positive probability. Part (3) is a continuity assumption that we use in the proof that an equilibrium exists.

Myerson [1998] showed that in games with stochastic participation, such as the spot-market in our model, beliefs over the total number of competitors from the perspective of a participant in the auction (i.e., a bidder) are not the same as beliefs from the perspective of an outside observer (e.g., a seller or the platform designer). The core intuition behind this phenomenon is that bidders make inferences about the size of an auction from the fact that he or she has been matched into that auction. For example, if π has heavy tails (which the estimates of Section 4 imply), then the distribution of auction sizes has more small and large auctions than a standard Poisson distribution would predict. Conditional on being matched into an auction, a bidder believes that it is likely that she is assigned to a large auction. From the perspective of bidder 1, let M denote a random variable representing the number of competitors she faces, and let $\pi_M(M; \boldsymbol{\lambda})$ denote its probability mass function

(PMF). Her beliefs over M are given by

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

To complete the description of the game, we need to describe the matching of bidders to auctions when the measure of entering bidders vanishes.[8] In this "no participation by buyers" limit (i.e., $\mathcal{C} \to 0$), each bidder that does choose to enter is the sole bidder in the auction with arbitrarily high probability.

Assumption 2.2. $\pi_M(0;\lambda) \to 1$ *as* $\mathcal{C} \to 0$*.*

Finally, we note at this point that our continuum model is "large" in the sense that the actions of individual bidders have no effect on the aggregate distribution of auction outcomes. However, the actions of individual bidders have a large effect on the auction to which they have been assigned.

2.2. **Equilibrium.** Because the model is stationary, a bidding strategy can be written as a function $\beta : [0,1] \to [0,1]$ with a typical bid denoted $\beta(v)$. The entry decision is a function $\zeta : [0,1] \to \{0,1\}$ where $\zeta(v) = 1$ indicates the agent enters the market and $\zeta(v) = 0$ that the agent stays out of the market. $\sigma = (\beta, \zeta)$ denotes the full strategy.

We use the notation $x(b, C, F_V, G_R|\sigma) = 1$ (0) to denote the random event that a buyer wins (loses) an auction with a bid of b given the aggregate state (C, F_V, G_R) and a common strategy vector for the buyers σ. $p(b, C, F_V, G_R|\sigma)$ denotes the random transfer from the buyer to the seller conditional on a bid of b given (C, F_V, G_R) and σ.[9] We assume throughout that $p(b, C, F_V, G_R) \leq b$. Both x and p are defined by the auction mechanism used on the platform. Where confusion will not result, we suppress the notation for the aggregate state variables C, F_V, and G_R and the strategy σ.

To simplify notation we also define

$$\begin{aligned} \chi(b, C, F_V G_R|\sigma) &= E\left[x(b, C, F_V, G_R|\sigma)\right] \\ \rho(b, C, F_V, G_R|\sigma) &= E\left[p(b, C, F_V, G_R|\sigma)\right] \end{aligned}$$

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. Note that ρ represents expected transfers that are not conditional on sale. That is, each entering bidder has an ex-ante expectation of paying $\rho(b)$ in the spot-market, although under any winner-pay pricing rule only one bidder will pay a positive amount, ex-post.

[8]Such an assumption is necessary to evaluate deviations by a buyer from a candidate equilibrium wherein no buyers enter the market.

[9]For example, $p(b, C, F_V, G_R|\sigma) = 0$ if the buyer does not win the auction.

All agents discount future payoffs using a per-period discount factor $\delta \in (0,1)$. The value function given a (symmetric) strategy vector $\sigma = (\beta, \zeta)$ played by all agents is:

$$\begin{aligned}\mathcal{V}(v, C, F_V, G_R|\sigma) =& \zeta(v)\left[\chi[\beta(v)]v - \rho[\beta(v)] - \kappa + (1-\chi[\beta(v)])\delta\mathcal{V}(v, C, F_V, G_R|\sigma)\right] \quad (2)\\ &+ (1-\zeta(v))\delta\mathcal{V}(v, C, F_V, G_R|\sigma)\\ =& \zeta(v)\left[\chi[\beta(v)](v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma)) - \rho[\beta(v)] - \kappa)\right] + \quad (3)\\ &\delta\mathcal{V}(v, C, F_V, G_R|\sigma)\end{aligned}$$

We use the notation $\mathcal{V}(v, C, F_V, G_R|\sigma_i', \sigma_{-i})$ when buyer i uses strategy σ' and all other agents follow σ. When describing agent behavior, we refer to "agents best-responding to the aggregate state," which we think better captures the economic intuition than the more conventional vernacular of "agents best responding to the actions of the other players." This is analogous to describing an agent in a general equilibrium economy as best responding to prices rather than the actions of the other agents that generate the prices.

Suppose that the distribution of types F_V and the strategy σ generate a distribution of bids G_B in the auctions on the platform. We let β^s denote the best-response in the static version (i.e., when $\delta = 0$) of the spot-market mechanism given a distribution of bids G_B and starting prices G_R. One useful property of our model is that we can describe the best-response to the aggregate state variables in terms of β^s. To see this, we first define a bidder's private value minus her opportunity cost as her *dynamic value*, denoted $\tilde{v}_v \equiv v - \delta\mathcal{V}(v)$. If a buyer enters the market, her optimal bid is defined by:

$$\underset{b}{argmax}\ \chi[b](v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma)) - \rho[b] = \underset{b}{argmax}\ \chi[b]\tilde{v}_v - \rho[b] \quad (4)$$

But this is the problem faced by a buyer with value $\tilde{v}_v$ in the static form of the spot-market. Thus, $\beta(v) = \beta^s(v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma))$ is a best-response bidding strategy to the aggregate state (C, F_V, G_R) and a strategy vector (β, ζ) played by the other buyers.

In most static spot-market mechanisms, one's equilibrium bid is chosen to balance out opposing forces: a higher bid will increase the chance of winning, but it may also raise the price one pays as well. Whenever the second force is present, bidders shade their demand. Henceforth, we refer to these two forces as *static incentives*. Intertemporal dynamics introduce an additional incentive for bidders to shade their bids: if the spot-market price is sufficiently high today, then a bidder would prefer to wait in expectation of lower prices tomorrow. Therefore, even when the spot-market game follows a second-price rule, rational bidders in equilibrium engage in demand shading. In what follows, we refer to this source of demand shading as *dynamic incentives*. Since a buyer with value e has a continuation value of 0, this means a bidder with a value of e is the only type of bidder that does not shade her bid (i.e., $\beta(e) = e$).

Proposition 2.3 summarizes four useful properties of the best-response, $(\widetilde{\beta}, \widetilde{\zeta})$, to (C, F_V, G_R).

Proposition 2.3. *The best-response strategy* $(\widetilde{\beta}, \widetilde{\zeta})$ *satisfies:*

(1) $\widetilde{\beta}(v) = \beta^s(v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma))$ *is a best-response bidding strategy.*
(2) $\widetilde{\beta}(v)$ *is increasing in* v *if* $\chi(b)$ *is increasing and uniquely defined for a set of* v *of Lebesgue measure 1.*[10]
(3) $\widetilde{\beta}(v)$ *if* $\chi(b)$ *is increasing.*
(4) *There exists a cutoff* $\widetilde{e}$ *such that* $\widetilde{\zeta}(v) = 1$ *if and only if* $v \geq \widetilde{e}$.

Part (1) summarizes our discussion of the relationship between the static and dynamic auctions. Part (2) uses the supermodularity of the buyer's decision problem to show that best-responses by the bidders must be monotonically increasing, and the monotonicity implies that almost all buyer-types have a unique best response. This insight also implies that the best-response function is symmetric across agents. Part (3) proves that if $\chi(b)$ is strictly increasing, then the best-response function is strictly increasing. Part (4) implies that we can describe the best-response strategy as $(\widetilde{\beta}, \widetilde{e})$, where $\widetilde{\zeta}$ has been replaced with the cutoff $\widetilde{e}$. With these results in hand, we can define our notion of stationary equilibrium.

Definition 2.4. The strategy vector $\sigma = (\beta, e)$ and the states C and $F_V \in \Delta([0,1])$ are a *Stationary Competitive Equilibrium* (SCE) if for all bidder values v we have:

(1) **Optimal Bids**: $\beta(v) = \underset{b}{argmax}\ \chi[b](v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma)) - \rho[b]$
(2) **Optimal Entry**: The entry cutoff is determined by the equation:
$$\chi[\beta(e)]e - \rho[\beta(e)] = \kappa \tag{5}$$
(3) **Stationarity**: The measure and distribution of agents entering the market equals those of the exiting agents:
$$\text{For all } v \geq e,\ \mu t_V(v) = \chi(\beta(v)) f_V(v)\mathcal{C} \tag{6}$$

Part (1) of Definition 2.4 implies the bidding strategy of buyers is optimal given that the buyer enters. Part (2) defines the indifference condition of buyers with the valuation that defines e. The lowest value of buyer that chooses to enter must be indifferent between entering and staying out. Since the equilibrium is stationary, buyers with a value $v < e$ prefer to stay out in every period (and hence buyers with these values exit the game immediately), while buyers with a value $v > e$ strictly prefer to enter in every period. The marginal buyer, $v = e$, is indifferent between entering and staying out and earning a payoff of 0. Therefore, the marginal agent's utility in any given period is:
$$\mathcal{V}(e, C, F_V, G_R|\sigma) = \chi[\beta(e)]e - \rho[\beta(v)] - \kappa = 0 \tag{7}$$
which yields the equation in part (2) of Definition 2.4.

[10] $\chi(b)$ is strictly increasing in our application since G_R has full support.

Part (3) of Definition 2.4 encapsulates what it means for the state variables C and F_V to be consistent with the laws of motion of the game. For an economy to be stationary, the distribution and measure of buyers that win auctions and exit the game must be replaced by an identical distribution of new entrants. Recall that μ is the measure of buyers entering each period with distribution t_V, so the left-hand side of Equation 6 describes the mass of agents of each type entering the game each period. $\chi(\beta(v))$ describes the probability an agent of type v wins and exits the economy, so $\chi(\beta(v))f_V(v)\mathcal{C}$ is the mass of agents of type v who win an auction and exit the game each period.[11]

To get some insight into the forces driving agent behavior in our model, consider the SPA spot-market mechanism. Since it is an equilibrium in weakly undominated strategies for a bidder to bid her value in the static version of this mechanism, the SCE strategy can be described as $\beta(v) = v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma)$.[12] In the static, one-shot setting, the opportunity cost (i.e., $\mathcal{V}(v, C, F_V, G_R|\sigma)$) is 0 as 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.

For our equilibrium existence and large-market approximation results (Appendix B) to hold, we require the following assumption on the equilibrium strategy of the spot-market mechanism. It is stated in terms of the spot-market equilibrium bidding strategy since the properties of these equilibria have been extensively characterized in the past.

Assumption 2.5. *Let $\mathcal{Q}[0,\overline{q}]$ denote the space of measures over $[0,1]$ that admit pdfs bounded from above by $0 < \overline{q} < \infty$. For any $\overline{q}$, the best response bidding strategy of the spot-market mechanism, $\widetilde{\beta}^s$, to G_B is continuous with respect to any $G_B \in \mathcal{Q}[0,\overline{q}]$; any distribution of starting prices, G_R, that has full-support with atoms only at 0; and the matching parameter, $\boldsymbol{\lambda}$.*

We can choose $\varphi \in (0,1)$ such that for any static best response $\widetilde{\beta}^s$ to $(\boldsymbol{\lambda}, G_B, G_R)$ where $G_B \in \mathcal{Q}[0,\overline{q}], \overline{q} < \infty$, and any $v > v'$ we have:

$$\widetilde{\beta}^s(v) - \widetilde{\beta}^s(v') \in \left[\varphi(v - v'), \frac{v - v'}{\varphi}\right] \tag{8}$$

It is difficult to provide underlying conditions that insure that Assumption 2.5 is satisfied. In particular cases, such as the SPA, it obviously holds. Since we have stated the assumption in terms of the equilibrium of the spot-market mechanism, we can use results on static auction to characterize $\widetilde{\beta}^s$ and check whether the assumption holds. For example, one can

[11]Another way of interpreting μ is that each spot-market auction has $1 - \mu$ probability of resulting in no sale because $k = 0$ players were matched to it or because the reserve price was too high.

[12]Similar equations appear in Jofre-Bonet and Pesendorfer [2003], Backus and Lewis [2016], Iyer, Johari, and Sundararajan [2014]. Jofre-Bonet and Pesendorfer [2003] studies a first-price procurement auction, so the formula includes a static demand-shading incentive. The latter two study SPA mechanisms that admit only dynamic bid-shading incentives.

show the assumption holds for a first-price auction by examining the differential equation describing the first order condition.[13]

Assumption 2.5 provides continuity properties that are essential for our proof of equilibrium existence. Our proof relies on a fixed point argument that requires the best-responses of the buyers and the market aggregates be continuous with respect to each other, which is the crucial contribution of the first half of Assumption 2.5. The remainder of the assumption mandates bounds on the slope of the bidding function. The lower bound insures that there will not be atoms in the bid distribution, which would destroy the model's continuity. The upper bound insures the set of bidding strategies we need to consider is Lipschitz continuous (and hence equicontinuous and compact). Together continuity and compactness allow us to apply Schauder's fixed point theorem to prove an equilibrium exists.

Proposition 2.6. *A stationary competitive equilibrium exists, and a positive mass of buyers choose to enter the market if κ is not too large.*

The main difficulties in the proof are (1) arguing we can limit consideration to a compact strategy and state space and (2) ruling out a number of utility discontinuities that naturally arise in auction markets. Once we handle these issues, our proof uses a traditional fixed point argument. We discuss the issue of uniqueness of the equilibrium in Section 5.

Finally, one must acknowledge that the model we employ uses a continuum of agents, whereas the actual eBay market is only used by a finite number of buyers and sellers in reality. Ideally one would have estimated a finite model that allows bidders to influence the aggregate state. However, the size of the state-space of such a finite model grows exponentially in the number of agents, and this in turn makes it impossible to solve the model and generate the counterfactuals we are interested in studying.

In Appendix B we prove that the stationary equilibrium we study is an approximate equilibrium of a finite model with sufficiently many agents. By an "approximate equilibirum," we mean that any individual agent has at most a small incentive to deviate from the stationary strategy when all agents play the SCE strategy. The intuition is that the stationary state of the continuum model is with high probability a good approximation of the aggregate state of a finite model with many agents. Since the bidding strategy is optimal with respect to the stationary state, the same strategy will be approximately optimal with respect to the actual state realized in the finite model. If there is little to be gained by costly monitoring of and optimizing with respect to the aggregate state realized in the finite-agent model, then bidders would find it optimal to act "as if" the aggregate state is fixed.

Our approximation result is more than just a theoretical nicety. Proving such a result requires us to state concretely the finite-agent market structures that we believe our model

[13]The FPA case relies on the assumption that G_R has full-support. If G_R did not possess full-support, then the bid function could become flat for high-value agents if G_B has a support limited to low bids.

approximates. In our case, we assume that as the market grows, a large number of auctions occur each day and the agents can only participate in a randomly chosen auction each day. In addition, approximation results typically rely on the continuity of the model, and our auction setting includes a large number of potential discontinuities that make it non-obvious that such an approximation result holds.

3. AN EMPIRICAL MODEL OF DYNAMIC PLATFORM MARKET BIDDING

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

For simplicity of discussion, consider the decision problem of a bidder who has decided to enter and finds herself competing within a spot-market auction; we will refer to her as bidder 1. As before, denote the total number of opponents she faces by $M \equiv K-1 \geq 0$ and recall that from 1's perspective $\pi_M(\cdot)$ may not be the same distribution as $\pi(\cdot)$, though $\boldsymbol{\lambda}$ determines both. Prior to bidding, bidder 1 observes her own private valuation v and she views her opponents' private values as independent realizations of a random variable $V \sim F_V$ having strictly positive density f_V on support $[\underline{v}, \overline{v}]$ with $\underline{v} \geq 0$.[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 Equation 5 determining the relevant entry cutoff. Since we are unable to collect real-world observations on non-entrants, we adopt the convention that F_V is the steady-state distribution of buyer types who choose *Enter*. Let $\underline{v} = e$ denote the type that is just indifferent to entering, leading to the following formula that we refer to as the "zero surplus condition":

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

Bidder 1 wishes to formulate an optimal bid that reflects her static incentives (from competition in the spot-market) and her dynamic incentives (from the option value of re-entering the market in the future if she loses today). She views the bids of her opponents as a random variable $B = \beta(V)$, which follows distribution $G_B(B) = F_V\left[\beta^{-1}(B)\right]$ with support $[\underline{b}, \overline{b}]$. Let B_M denote the maximal bid among all of bidder 1's opponents, and we adopt the convention that $B_M = 0$ if there are no other opponent bidders. The distribution of B_M is then $G_{B_M}(B_M) = \pi_M(0) + \sum_{m=1}^{\infty} \pi_M(m) G_B(B_M)^m$. R is the starting price of the auction,

[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 compact interval of real numbers.

which is randomly drawn from CDF G_R. In order to win, player 1's bid must exceed the realized value of the random variable $Z \equiv \max\{R, B_M\}$. Note that the distribution of Z is the same as the win probability function:

$$\chi(b) = G_R(b) \sum_{m=0}^{\infty} \pi_M(m) G_B(b)^m = G_Z(b) \tag{9}$$

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

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

$$\rho(b) \equiv \mathrm{E}[p_B(b)] = \int_0^b t g_Z(t) dt \tag{10}$$

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

$$\mathcal{V}(v) = \max_{b \in \mathbb{R}_+} \left\{ \chi(b) v - \rho(b) - \kappa + [1 - \chi(b)] \, \delta \mathcal{V}(v) \right\} \tag{11}$$

$$\beta(v) = v - \delta \mathcal{V}(v) \tag{12}$$

The demand shading factor given by bidder 1's continuation value, $\delta\mathcal{V}(v)$, is uniquely characterized by four things: the per-period entry cost, κ; the distribution of bids, $G_B(b)$; the distribution of starting prices, G_R; and the market tightness parameters, $\boldsymbol{\lambda}$, that determine the overall ratio of buyers and sellers. Thus, mapping bids into private values requires first identifying these four objects. The model is said to be identified if there exists a unique set of structural primitives that could rationalize a given realization of the joint distribution of observables, $\{\tilde{k}_l, r_l, y_l\}_{l=1}^{L}$. The structural primitives to identify are μ, κ, T_V, and G_R.[15]

Note that 12 implies $\mathcal{V}(v) = \frac{v - \beta(v)}{\delta}$. By substituting this expression into equation 11 and using the shorthand notation $b^* = \beta(v)$, we can rearrange terms to get

$$v = b^* \frac{1 - \delta\,(1 - \chi(b^*))}{1 - \delta} - \frac{\delta}{1 - \delta} \Big(\rho(b^*) + \kappa \Big) \equiv \beta^{-1}(b^*) \tag{13}$$

[15]Our exposition might appear to treat F_V as a primitive, but in reality it is endogenously pinned down by T_V (and other variables) through equation 6. Conversely, after estimating the primitive variable μ and the endogenous variables F_V and C, we can recover the primitive T_V. Our counterfactuals hold T_V (the distribution of potential entrant types) fixed and allow F_V (the steady-state distribution of active bidder types) to vary endogenously with alternate market structures.

The expression for the inverse bidding function above is crucial to demonstrating identification of the model, as it shows 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 bid b as a best response to prevailing market conditions.

3.1.2. *Identification of the Bid Distribution and Bidder Arrival Process.* Proofs of our main identification results in this section are relegated to Appendix A.3. If the econometrician knew exactly how many bidders attended each auction, then the empirical problem here would be a simple one to solve. For example, if the number of bidders was known to be k, then the CDF of the (observable) highest losing bid, call it $H_k(b)$, is related to its (unobserved) parent distribution through the following bijective mapping:

$$\begin{aligned} H_k(b) &= G_B(b)^k + kG_B(b)^{k-1}\left[1 - G_B(b)\right] \\ &= \binom{k}{k-2} \int_0^{G_B(b)} u^{k-2}(1-u)^k du \equiv \phi_k^{-1}\left(G_B(b)\right). \end{aligned}$$

From the above integral it is easy to see that ϕ_k is a monotone bijection that maps quantile ranks of $H_k(b)$ into quantile ranks of $G_B(b)$: the upper limit of integration ranges from the lower limit to one, and since the integrand is always positive, it is monotone in G_B as well. Our setting is more complex though: since the the number of bidders is unknown and random, the observed distribution of the highest losing bid, $H(\cdot)$, takes the form

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

In other words, $H(b)$ 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. Note that since $\phi^{-1}\left(G_B(b)\right)$ is simply a convex combination of monotone bijective transformations $\phi_k^{-1}\left(G_B(b)\right)$, it must also be a bijection as well.

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

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

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

matched bidders, this filter process first randomly assigns each bidder an index $\{1, 2, \ldots, k\}$, where one's position in the list determines the ordering of bid submission times. Nature then visits each bidder in the order of her index within the list, keeping a running record of the current lead bidder and current price as she goes. As Nature visits each bidder in the list, she only records bid tenders that cause her running record of the price or lead bidder to update (i.e., those that exceed the second highest from among previous bid tenders). Otherwise, Nature skips bidder i's submission as if it never happened and reports to the econometrician only the record of price path updates, which reveals $\tilde{k} \leq k$ observed bidder identities. This filter process is meant to depict the way information is recorded on real-world platform markets like eBay, and it opens up the possibility that some bidders will not appear to have participated even though they had an intent to bid. This view of intra-auction dynamics assumes that the ordering of bidders' submission times is random. Note that we remain agnostic on how agents decide to bid early or late; we merely rule out the possibility of bidder collusion on the ordering of their submissions.

Since the filter process does not depend on the particular distribution of K, the distribution of $\tilde{K}$ conditional on a given k can be characterized without knowing $\boldsymbol{\lambda}$. Moreover, since a bidder's visibility to the econometrician only depends on whether her bid exceeds the second-highest preceding bid, the researcher can easily simulate the filter process based on quantile ranks—without knowing G_B or F_V a priori—to compute conditional probabilities $\Pr[\tilde{k}|k]$, the probability of observing random lower bound $\tilde{k}$ given that k total bidders were actually competing, for various $(\tilde{k}, k)$ pairs.[17] We adopt a special notation for this function, $P_0(\tilde{k}, k) \equiv \Pr[\tilde{k}|k]$ and treat it as known since it can be computed without data on bidding. 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 the presence of binding reserve prices. These introduce a second layer of selection, driving a further wedge between actual participation k and observed participation, $\tilde{k}$. Not only do some bidders go unobserved because the filter process withholds them from view, but an additional fraction of bidders, who would have otherwise been reported by Nature, go unobserved because their bids fall below the reserve price. This second layer of selection produces substantial complications since G_B now determines how the second source of selection influences the relation between the distribution of observed $\tilde{K}$ and the underlying distribution of actual K.

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

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

In order to characterize the conditional distribution of $\tilde{K}$ given r, first note that if there are $K = k$ total bidders, the probability that exactly j of them are screened out by r is $\binom{k}{j} G_B(r)^j \left[1 - G_B(r)\right]^{k-j}$. Now suppose there are $\tilde{K}$ observed bidders in an auction with K total bidders. We can combine the two levels of selection in the adjusted filter process with the following equation:

$$\Pr[\tilde{K} = \tilde{k} | K = k, r] = \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} G_B(r)^j \left[1 - G_B(r)\right]^{k-j} P_0(\tilde{k}, k - j). \tag{15}$$

The sum is to account for the fact that any number of bidders between 0 and $k - \tilde{k}$ could be screened out by selection on reserve prices. The trailing term on the end involving P_0 is to account for the standard filter process running its course with the surviving set of bidders. Equation (15) now allows us to characterize the distribution of observed $\tilde{K}$ conditional on the observable reserve price r, as

$$\tilde{\pi}(\tilde{k}|r) = \sum_{k=\tilde{k}}^{\infty} \Pr[\tilde{k}|k, r] \pi(k; \boldsymbol{\lambda}). \tag{16}$$

Estimation of $\boldsymbol{\lambda}$ can no longer be separated from G_B because Equation 16 involves both of these objects. Fortunately though, this is merely a matter of implementation, as the following two results demonstrate sufficient conditions under which the model is nonparametrically identified from the available observables. We begin by demonstrating identification for a finite-dimensional restriction of the model:

Lemma 3.2. *Suppose that the market tightness parameters take the form $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\}$, where $\lambda_k \equiv \pi(k; \boldsymbol{\lambda})$ for each $k = 0, 1, 2, \ldots$ Suppose further that all but I of them are zero, so that we can re-express them as $\boldsymbol{\lambda} = \{\lambda_{k_1}, \lambda_{k_2}, \ldots, \lambda_{k_I}\}$, where $k_1 < k_2 < \cdots < k_I$. Then it follows that there is a unique $(\boldsymbol{\lambda}, G_B)$ pair that is consistent with the joint distribution of the observables $(\tilde{K}, R, Y)$ for any finite I.*

The proof of the lemma (see Appendix A.3) does not require any shape restrictions on G_B, and it does not depend on the specific value of I in any way (aside from its finiteness). In other words, the identification result holds for models of the bidder arrival process that become arbitrarily flexible as I grows. In that sense, the preceeding result may be considered as nonparametric.

The proof of the lemma is non-constructive in the sense that it does not immediately suggest a practical way of obtaining parameter estimates from a finite sample of data. For a more intuitive understanding of how the observables pin down a unique set of parameters under the model, consider that bijectivity of the mapping ϕ implies that fixing a value for $\boldsymbol{\lambda}$ pins down a unique $G_B^{\lambda}(b) = \phi\left(H(b); \boldsymbol{\lambda}\right)$ through equation (14). Therefore, if we condense equations (14)–(16) we can write

$$\tilde{\pi}(\tilde{k}|r) = \sum_{k=\tilde{k}}^{\infty} \pi(k; \boldsymbol{\lambda}) \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} \phi\left(H(b); \boldsymbol{\lambda}\right)^{j} \left[1 - \phi\left(H(b); \boldsymbol{\lambda}\right)\right]^{k-j} P_0(\tilde{k}, k-j), \tag{17}$$

which depends only on $\boldsymbol{\lambda}$ and the observables. In principle then, for any finite, I-dimensional, parametric restriction of the bidder arrival process $\pi(k; \boldsymbol{\lambda})$, equation (17) allows the researcher to choose from a continuum of moment conditions, with each possible moment condition depending on a distinct $(\tilde{k}, r)$ pair. Using Lemma 3.2 we are now ready to state the first of our main identification results.

Proposition 3.3. *For a given discount factor δ, suppose that the market tightness parameters take the form $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\}$, where $\lambda_k \equiv \pi(k; \boldsymbol{\lambda})$ for each $k = 0, 1, 2, \ldots$ Suppose further that all but I of them are zero, so that we can re-express them as $\boldsymbol{\lambda} = \{\lambda_{k_1}, \lambda_{k_2}, \ldots, \lambda_{k_I}\}$, where $k_1 < k_2 < \cdots < k_I$. Then when the spot-market mechanism is a sealed-bid, second-price auction, for any finite I there is a unique configuration of model parameters $\boldsymbol{\Theta} \equiv (\boldsymbol{\lambda}, \kappa, F_V, \mu, T_V)$ that is consistent with the joint distribution of the observables $\left(\tilde{K}, R, Y\right)$.*

Once again, this result may be considered as nonparametric in the sense that it requires no shape restrictions on G_B, F_V or T_V, and it holds for models of the bidder arrival process that become arbitrarily flexible as I grows.

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

Recall that a buyer with valuation v has a dynamic value equal to $\widetilde{v}_v = v - \delta\mathcal{V}(v)$. Since a buyer bids as if they are playing a static auction with their dynamic bid shading incentives collapsed into their dynamic value (Proposition 2.3), we can show that if the right set of observables are available to identify the mapping β^s 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 β^s into equation (11) and rearranging we get:

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

Using the shorthand $b^* = \beta^s(\tilde{v}_v) = \beta(v)$ 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 SPA spot-market, where $b^* = \beta^s(\tilde{v}_v) = \tilde{v}$, equation (18) reduces to equation (13) above.

Proposition 3.4. *For a given discount factor δ, suppose that the market tightness parameters take the form $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\}$, where $\lambda_k \equiv \pi(k;\boldsymbol{\lambda})$ for each $k = 0, 1, 2, \ldots$ Suppose further that all but I of them are zero, so that we can re-express them as $\boldsymbol{\lambda} = \{\lambda_{k_1}, \lambda_{k_2}, \ldots, \lambda_{k_I}\}$, where $k_1 < k_2 < \cdots < k_I$. Then for any finite I there is a unique configuration of model parameters $\boldsymbol{\Theta} \equiv (\boldsymbol{\lambda}, \kappa, F_V, \mu, T_V)$ that is consistent with the joint distribution of the observables $\big(\tilde{K}, R, Y\big)$ under any spot-market mechanism for which the optimizer of equation 4, $\beta^s(v)$, could be identified from the available observables $\big(\tilde{K}, R, Y\big)$ if they were generated from a sample of static, one-shot auction games.*[18]

Proof. The argument for $\boldsymbol{\lambda}$, G_B, and G_R is the same as in Lemma 3.2. Now consider a hypothetical 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, $\beta^s(\cdot)$, 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 equation (6) and the steady-state identity $\mu = \int_{\underline{v}}^{\overline{v}} \chi\left[\beta(v)\right] f_V(v)dv$ identify μ and T_V. □

Proposition 3.4 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

[18]Alternatively, if the optimizer of equation (4) is unique and the allocation rule $\chi(b)$ and pricing rule $\rho(b)$ can be identified from the available observables $\big(\tilde{K}, R, Y\big)$, then $\beta^s(v)$ is identified.

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

3.2. **A Two-Stage, Semi-Parametric Estimator.** Thus far in our discussion we have left the bidder arrival process $\pi(k;\boldsymbol{\lambda})$ 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 now assume K follows a generalized Poisson distribution (Consul and Jain [1973]) with PMF:

$$\pi(K=k;\boldsymbol{\lambda}) = Pr\left[K=k|\boldsymbol{\lambda}\right] = \lambda_1(\lambda_1+k\lambda_2)^{k-1}\frac{e^{-(\lambda_1+k\lambda_2)}}{k!}, \;\; \lambda_1>0, \;\; |\lambda_2|<1. \tag{19}$$

The first two moments of the generalized Poisson distribution are $\mathrm{E}[K] = \lambda_1/(1-\lambda_2)$ and $\mathrm{Var}[K] = \mathrm{E}[K]/(1-\lambda_2)^2$. While the generalized Poisson reduces to a regular Poisson distribution when $\lambda_2 = 0$, it exhibits fatter tails when $\lambda_2 > 0$ and thinner tails when $\lambda_2 < 0$. Given the linkage between the traditional and generalized Poisson distributions, we refer to λ_1 as the *size parameter* and λ_2 as the *dispersion parameter*.

Recall 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.

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. 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 $[\hat{\underline{b}}, \hat{\bar{b}}] = [\min_l\{y_l\}, \max_l\{y_l\}]$ 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 $n_{b,I_b} = \hat{\bar{b}}$ 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 $\mathcal{F}_{bi} : [\hat{\underline{b}}, \hat{\bar{b}}] \to \mathbb{R}$, $i = 1, \ldots, I_b + 3$ that give us our parameterization of the bid distribution:[19]

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

We also follow this approach for estimating G_R and F_V. Let $\mathbf{n}_r = \{n_{r1} < n_{r2} < \cdots < n_{r,I_r+1}\}$ 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.[20] The former is chosen to span $[\underline{r}, \hat{\bar{r}}] = [0.99, \max_l\{r_l\}]$ and the latter spans $[\hat{\underline{v}}, \hat{\bar{v}}]$, with the bounds to be estimated. These knot vectors determine our other basis functions $\mathcal{F}_{r,i} : [\underline{r}, \hat{\bar{r}}] \to \mathbb{R}$, $i = 1, \ldots, I_r + 3$ and $\mathcal{F}_{v,i} : [\hat{\underline{v}}, \hat{\bar{v}}] \to \mathbb{R}$, $i = 1, \ldots, I_v + 3$ which in turn render our parameterizations $\hat{G}_R(r; \boldsymbol{\alpha}_r) = \sum_{i=1}^{I_r+3} \alpha_{ri} \mathcal{F}_{r,i}(r)$ and $\hat{F}_V(v; \boldsymbol{\alpha}_v) = \sum_{i=1}^{I_v+3} \alpha_{vi} \mathcal{F}_{v,i}(v)$.

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, $(\beta^s)^{-1}(\cdot)$, $\beta^{-1}(\cdot)$, $\mathcal{V}(\cdot)$, $F_V(\cdot)$, μ, and T_V as functions of first-stage parameter estimates. Note that Stages I and II differ in that Stage I is an estimation step, but Stage II is a computational step based on the outputs from Stage I.

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

$$\tilde{\pi}(\tilde{k}|r; \boldsymbol{\lambda}, \boldsymbol{\alpha}_b) = \sum_{k=\tilde{k}}^{\overline{K}} \left\{ \sum_{j=0}^{k-\tilde{k}} \binom{k}{j} \hat{G}_B(r; \boldsymbol{\alpha}_b)^j \left[1 - \hat{G}_B(r; \boldsymbol{\alpha}_b)\right]^{k-j} P_0(\tilde{k}, k-j) \right\} \pi(k; \boldsymbol{\lambda}) \tag{20}$$

[19] A standard text on B-splines is de Boor [2001]. See also [Hickman et al., 2016, Online Appendix].

[20] For discussion on choice of knot locations used in our empirical implementation, see Appendix C.1.

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:

$$\hat{\tilde{\pi}}(\tilde{k}|r) = \sum_{l=1}^{L} \mathbb{1}(\tilde{k}_l = \tilde{k}) \frac{\mathcal{K}\left(\frac{r-r_l}{h_R}\right)}{\sum_{t=1}^{L} \mathcal{K}\left(\frac{r-r_t}{h_R}\right)} \tag{21}$$

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

$$H(b; \boldsymbol{\lambda}, \boldsymbol{\alpha}_b) = \sum_{k=2}^{\infty} \frac{\pi(k; \boldsymbol{\lambda}) \left(G_B(b; \boldsymbol{\alpha}_b)^k + kG_B(b; \boldsymbol{\alpha}_b)^{k-1} \left[1 - G_B(b; \boldsymbol{\alpha}_b)\right]\right)}{1 - \pi(0; \boldsymbol{\lambda}) - \pi(1; \boldsymbol{\lambda})}, \tag{22}$$

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

$$\begin{aligned} (\hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b) = \underset{(\boldsymbol{\lambda}, \boldsymbol{\alpha}_b) \in \mathbb{R}^{I_b+5}}{\arg\min} & \sum_{l=1}^{L} \left\{ \left[\tilde{\pi}(\tilde{k}_l | r_l; \boldsymbol{\lambda}, \boldsymbol{\alpha}_b) - \hat{\tilde{\pi}}(\tilde{k}_l | r_l)\right]^2 + \left[H(y_l; \boldsymbol{\lambda}, \boldsymbol{\alpha}_b) - \hat{H}(y_l)\right]^2 \right\} \\ & \text{subject to} \\ & \alpha_{b,1} = 0, \quad \alpha_{b,I_b+3} = 1, \\ & \alpha_{b,i} \leq \alpha_{b,i+1}, \ i = 1, \ldots, I_b + 2. \end{aligned} \tag{23}$$

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 $\hat{G}_B$.

[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]An analogous estimator in the absence of the generalized Poisson assumption would be possible, but with additional complications. In the case where $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\}$ and $\lambda_k \equiv \Pr[K = k]$, the main challenge is that only finitely many elements of $\boldsymbol{\lambda}$ can be estimated with finite sample size L. Thus, one could choose an upper bound $\overline{K}_L < \infty$ and restrict $\lambda_k = 0$ for each $k > \overline{K}_L$. A fully nonparametric estimator must also specify the rate at which $\overline{K}_L$ should grow with the sample size. While interesting, the answers to these questions are beyond the scope of the current exercise, so we do not address them here. In a simpler setting than ours—a static bidding model of eBay laptop computer auctions with no binding reserve prices—Hickman et al. [2016] found strong evidence that the generalized Poisson assumption produced estimates of the bidder arrival process that could not be improved upon by relaxations of its parametric form, given their sample size of roughly 750 auctions.

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

$$\begin{aligned} \hat{\boldsymbol{\alpha}}_r = \arg\min_{\boldsymbol{\alpha}_r \in \mathbb{R}^{I_r+3}} & \sum_{l=1}^{L} \left\{ \left[\hat{G}_R(r_l; \boldsymbol{\alpha}_r) - \ddot{G}_R(r_l) \right]^2 \right\} \\ & \text{subject to} \\ & \alpha_{r1} = \ddot{G}_R(\underline{r}), \;\; \alpha_{r,I_r+3} = 1, \\ & \alpha_{ri} \leq \alpha_{r,i+1}, \; i = 1, \ldots, I_r + 2, \end{aligned} \tag{24}$$

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

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

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

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

$$v = b + \frac{G_Z(b) - G_Z\left[\tau(b)\right]}{g_Z(b)}, \;\; \tau(b) = \begin{cases} \underline{b} & \text{if } b \leq \underline{b} + \Delta \\ b - \Delta & \text{otherwise} \end{cases} \tag{25}$$

where $\tau(b)$ is a threshold function determining the point below one's own bid which, if the random variable Z surpasses it, will trigger a first-price outcome. Proposition 3.4 enables us to adapt equation (25) above for the static inverse bid function $(\beta^s)^{-1}$ in our model. Our inverse static bid function is given by:

$$\left(\hat{\beta}^{s}\right)^{-1}(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r) = \hat{\tilde{v}}_b = b + \frac{\hat{G}_Z(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r) - \hat{G}_Z\left[\tau(b);\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r\right]}{\hat{g}_Z(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)}. \tag{26}$$

Using Stage I estimates we can construct the allocation rule:

$$\chi(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r) = \hat{G}_R(b;\hat{\boldsymbol{\alpha}}_r)\sum_{m=0}^{\infty}\pi_M(m;\hat{\boldsymbol{\lambda}})\hat{G}_B(b;\hat{\boldsymbol{\alpha}}_b)^m. \tag{27}$$

Taking into account the hybrid pricing rule, we can construct the payment function:

$$\begin{aligned}\rho(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r) = \underline{r}G_R(\underline{r};\hat{\boldsymbol{\alpha}}_r) + \int_{\underline{r}}^{\tau(b)}(t+\Delta)\hat{g}_Z(t;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)dt \\ + b\left(\hat{G}_Z[b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r] - \hat{G}_Z[\tau(b);\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r]\right).\end{aligned} \tag{28}$$

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

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

$$\hat{\kappa} = \chi(\hat{\underline{v}}_{\underline{b}};\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)\hat{\underline{v}}_{\underline{b}} - \rho(\hat{\underline{v}}_{\underline{b}};\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r) \tag{29}$$

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

$$\hat{v} = \hat{\beta}^{-1}\left(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right) = \hat{\tilde{v}}_v\frac{1-\delta\left[1-\chi(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)\right]}{1-\delta} - \frac{\delta\left(\rho(b;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r)+\hat{\kappa}_B\right)}{1-\delta} \tag{30}$$

$$\hat{\mathcal{V}}\left(v;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right) = \frac{\hat{v}-\hat{\tilde{v}}_b}{\delta}. \tag{31}$$

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

$$\begin{aligned}\hat{\boldsymbol{\alpha}}_v = \underset{\boldsymbol{\alpha}_v\in\mathbb{R}^{I_v+3}}{\arg\min}\sum_{j=1}^{J}\left\{\left[\hat{G}_B\left(b_j;\hat{\boldsymbol{\alpha}}_b\right) - \sum_{i=1}^{I_v+3}\alpha_{vi}\mathcal{F}_{vi}\left[\hat{\beta}^{-1}\left(b_j;\hat{\boldsymbol{\lambda}},\hat{\boldsymbol{\alpha}}_b,\hat{\boldsymbol{\alpha}}_r,\delta\right)\right]\right]^2\right\} \\ \text{subject to} \\ \alpha_{v,1} = 0,\ \ \alpha_{v,I_v+3} = 1, \\ \alpha_{v,i} \le \alpha_{v,i+1},\ i = 1,\ldots,I_v+2.\end{aligned} \tag{32}$$

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

$$\hat{\mu} = \left[1 - \pi\left(0; \hat{\boldsymbol{\lambda}}\right)\right] G_R\left(\underline{b}; \hat{\boldsymbol{\alpha}}_r\right) + \int_{\underline{b}}^{\overline{b}} g_R\left(r; \hat{\boldsymbol{\alpha}}_r\right) \left(\sum_{k=1}^{\infty} \pi\left(k; \hat{\boldsymbol{\lambda}}\right) \left[1 - G_B\left(r; \hat{\boldsymbol{\alpha}}_b\right)^k\right]\right) dr \tag{33}$$

$$t_V\left(v; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b; \hat{\boldsymbol{\alpha}}_r, \delta\right) = \frac{\chi\left[\beta\left(v; \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b, \hat{\boldsymbol{\alpha}}_r, \delta\right); \hat{\boldsymbol{\lambda}}, \hat{\boldsymbol{\alpha}}_b, \hat{\boldsymbol{\alpha}}_r\right] f_V\left(v; \boldsymbol{\alpha}_v, \delta\right) \frac{\lambda_1}{1-\lambda_2}}{\hat{\mu}}. \tag{34}$$

3.2.3. *Asymptotics and Standard Errors.* Our estimators $\hat{\boldsymbol{\lambda}}$, $\hat{\boldsymbol{\alpha}}_b$, and $\hat{\boldsymbol{\alpha}}_r$ belong to a broad class of semiparametric M-estimators for which large-sample properties have been studied at length. In particular, since B-splines are mathematically equivalent to piecewise splines with differentiability conditions imposed at the interior knots (see de Boor [2001]), our Stage I estimators fit within a subset of this broad class–spline series least squares estimators–for which consistency and pointwise asymptotic normality are known (see Chen [2007] and Huang [2003]).[23] Since Stage II empirical objects are all smooth transformations of the functionals estimated in Stage I, they are also asymptotically pointwise normal. In our implementation we use the standard nonparametric bootstrap as a computationally convenient method of accounting for the role of sampling variability. Specifically, we re-sample from our auction-level observations (with replacement) to construct 1000 bootstrap samples of size L. Then we execute Stages I and II of our estimator as described above to obtain pointwise confidence bounds on functionals, as well as confidence bounds on a small number of interpretable parameter estimates.

4. DATA AND RESULTS

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

Each Kindle tablet had eight gigabytes of internal storage and a seven-inch screen with standard-definition resolution of 1024x600. The Kindle Fire tablets came pre-loaded with Amazon's proprietary version of the Android-based operating system that prevents the user from accessing the full Android app market.[24] This makes the Kindle Fire a poor substitute for a standard tablet (e.g., Samsung Galaxy or Apple iPad) that can serve a dual

[23]Here, pointwise normality refers to the estimated functionals from Stage I, $G_B(b; \hat{\boldsymbol{\alpha}}_b)$ and $G_R(r; \hat{\boldsymbol{\alpha}}_r)$, rather than to the parameters themselves.

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

role as a productivity tool or as a highly versatile consumer electronic device. Rather, the Kindle Fire is specifically designed to be a consumer access point exclusively to Amazon.com's electronic media market, which includes e-books, periodicals, audiobooks, music, and movies.[25] All transactions were covered by the eBay Money Back Guarantee to insure consumers against potential unscrupulous sellers.[26]

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

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

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

[25] 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.

[26] 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.

the UH variable, he shows that one can use joint movement in reserve prices and bids to deconvolve the UH and identify private valuations.

In our data we observe non-trivial variation in sellers' reserve prices with roughly one third of them being binding for a positive fraction of the bidder population. Therefore, one might reasonably suspect that if UH is present then higher values of the unobserved characteristic prompt sellers to increase reserve prices.[27] A necessary condition for UH in the Roberts model is co-movement of bids and reserve prices, which is testable. We find in our data that the correlation between seller reserve price and the highest losing bid is very small in a practical sense, at -0.015, and that it is also statistically indistinguishable from zero. We interpret the combined information above—lack of correlation between bids and reserve prices, evidence from our manual survey of `.html` pages, uniform buyer insurance, proprietary operating system and limited app market, and uniform characteristics of the Kindle Fire tablets—as evidence consistent with our assumption of a homogeneous goods market with no close substitutes. These characteristics of the eBay Kindle data allow us to avoid significant complications covered by other work, such as identifying UH or complex substitution patterns (see Backus and Lewis [2016]), and instead focus on questions of bidding behavior, allocative efficiency, and market design.

4.1. **Practical Concerns.**

4.1.1. *Intra-Auction Dynamics.* For each auction listing, we observe the timing and amount of each bid submission as well as an anonymized hash of 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 price approaches a reasonable level. Some bidders may engage in non-equilibrium cheap-talk before bidding based on best-response calculations or participate flippantly to pass time while web surfing. Empirically, a significant fraction of observed bid amounts, especially those submitted early in the life of the auction, fall too far below realistic transaction prices to be taken seriously. The question of intra-auction dynamics is broad, complicated, and beyond the scope of this work. In our case, inter-auction dynamics are the primary concern for answering our research questions on allocative efficiency and market design.

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

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

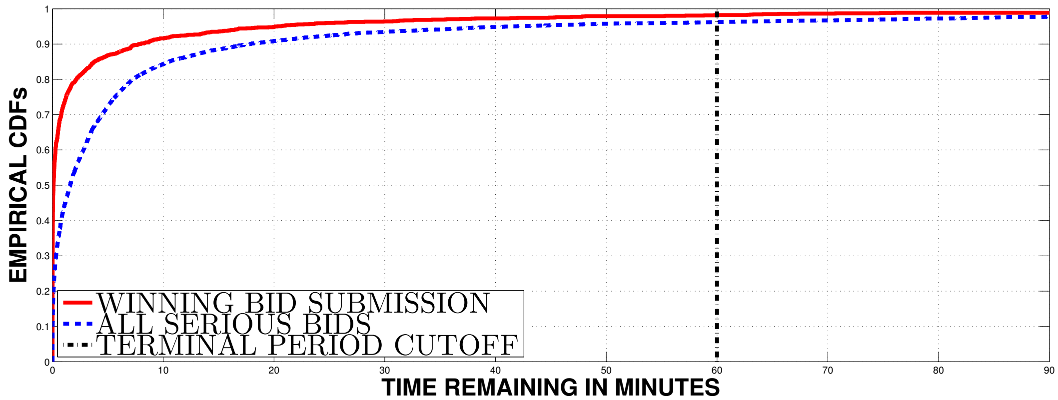

FIGURE 3. Empirical Distributions: Time Remaining when Bids are Submitted

bids and the final sale price. The second stage is treated as a sealed-bid auction as per our model of Sections 2 and 3.[28] 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. Note that our definition of serious bidding will also count the top two submissions from within the first stage of the auction as these bids fix the price at the start of the second stage of the auction. This allows us to avoid drawing too sharp a distinction between the cheap-talk stage and the terminal stage of the auction, since some serious bidders' submission times may still occur early in the life of the auction. Likewise, a *serious bidder* is one who is observed to submit at least one serious bid. Of course, the possibility always exists that some bidders who are determined to be non-serious by the above criterion had serious intent to compete for the item, but were priced out before submitting their planned, serious bid during the terminal stage. This is, however, part of the problem that our model of the adjusted filter process solves (i.e., observed participation by serious bidders is a lower bound on actual participation).

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

[28]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

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 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 only that the highest losing bid is fully reflective of equilibrium play. This leaves us with the three data points from each auction needed for identification: $\tilde{k}_l$, the observed number of serious bidders; r_l, the seller's reserve price; and y_l, the highest loser bid from auctions with at least two bids. After dropping `.html` pages for which our software was unable to extract data because of formatting problems, we have 1,462 total auctions, 2 of which logged no bids, and 1,397 of which had 2 or more observed bidders (so that we observed a highest losing bid). Table 1 displays descriptive statistics on bid timing, observed participation, sale prices, and highest losing bids.

4.2. **Choosing δ.** The final free parameter is the time discount factor, δ. As in many other empirical contexts, this part poses a difficult challenge. Luckily, δ does not enter Stage I estimation, so all of the necessary building blocks to compute the final structural primitives will be unaffected. Several Stage II objects are also unaffected, including the win probability $\chi(\cdot)$, the expected payment function $\rho(\cdot)$, the per-period participation cost $\hat{\kappa}$, and the exogenous, per-period measure of new agents flowing into the market $\hat{\mu}$. However, the remaining objects depend on δ. The objects affected by δ include the dynamic bid function

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)

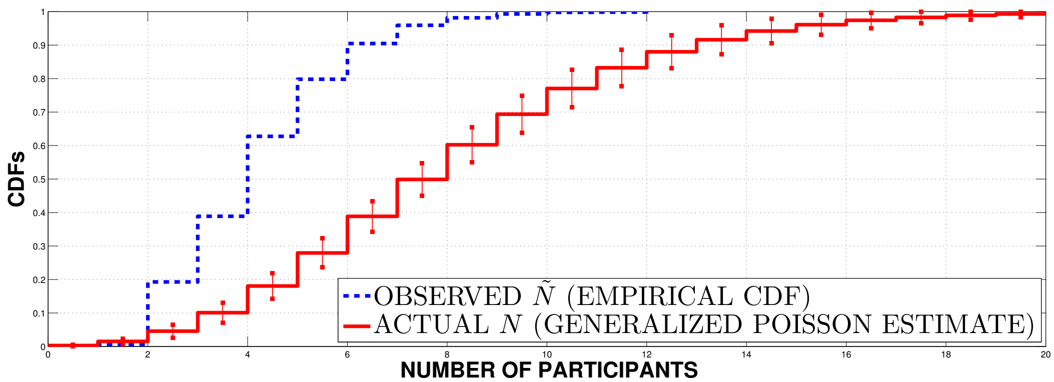

FIGURE 4. Matching Process Estimates

$\beta(\cdot)$, the value function $\mathcal{V}(\cdot)$, and the steady-state private value distributions for market participants $F_V(\cdot)$ and new entrants $T_V(\cdot)$. There is an intuitive reason why: these objects tell us something about the opportunity cost of losing today, and δ plays a pivotal role in shaping this opportunity cost by determining agents' attitude toward present versus future consumption.

In lieu of taking a stand on the particular value of δ applicable to our study, we present results both here and in our counterfactual section for a range of values of δ. Where possible, we provide statistics that are stable across choices of δ. For example, instead of providing a dollar value for deadweight loss, which is sensitive to δ, we present deadweight loss as a percentage of the buyer's value, which is stable across different choices of δ.

4.3. **Estimates.** Table 2 displays point estimates and standard errors for readily interpretable parameters, including the market tightness parameters, the per-period participation cost, and the per-period measure of new entering agents. Figure 4 depicts point estimates for the empirical CDF of observed bidders (thick, dashed line), the estimated distribution of total auction-level participation K (thick, solid line), and point-wise confidence bounds for a selected grid of domain points (vertical box plots). As the figure demonstrates, failing to account for unobserved bidders would lead to a very different view of the distribution of auction participation. This substantial difference shows up in both the mean—4.07 for observed bidders per auction versus 7.96 for actual bidders per auction—and also in the

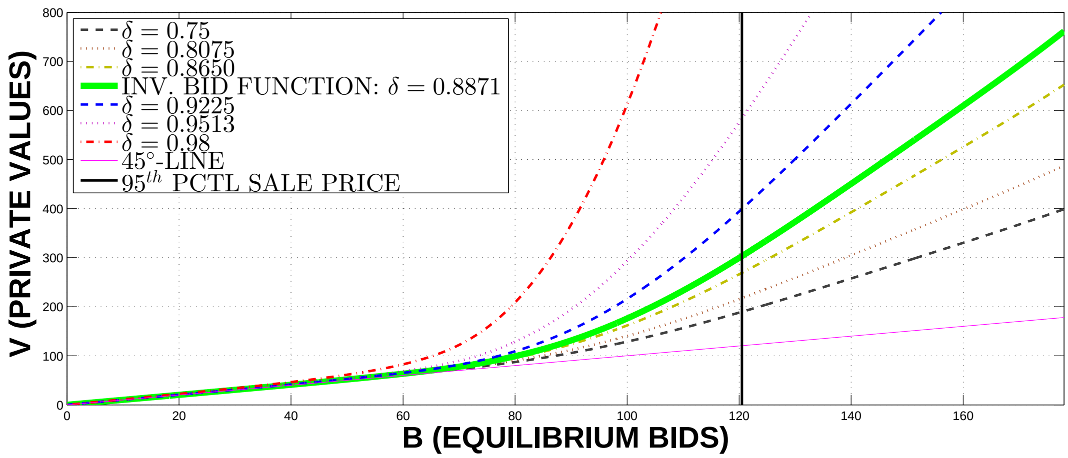

FIGURE 5. Inverse Bid Function Estimates Given Various Values of δ

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%

variance—3.19 for observed bidders and 14.46 for actual bidders. Appendix C.1 contains similar plots of point estimates for $\hat{G}_B(b; \boldsymbol{\alpha}_b)$ and $\hat{G}_R(r; \boldsymbol{\alpha}_r)$.

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 uniform grid of values of the time discount factor δ between 0.75 and 0.98. We also include an additional value at 0.8871 taken from an experimental study by Augenblick, Niederle, and Sprenger [2016] where they elicited hyperbolic time discount parameters at the daily level from college students.[29] 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. Continuation value is primarily determined by three things: the

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

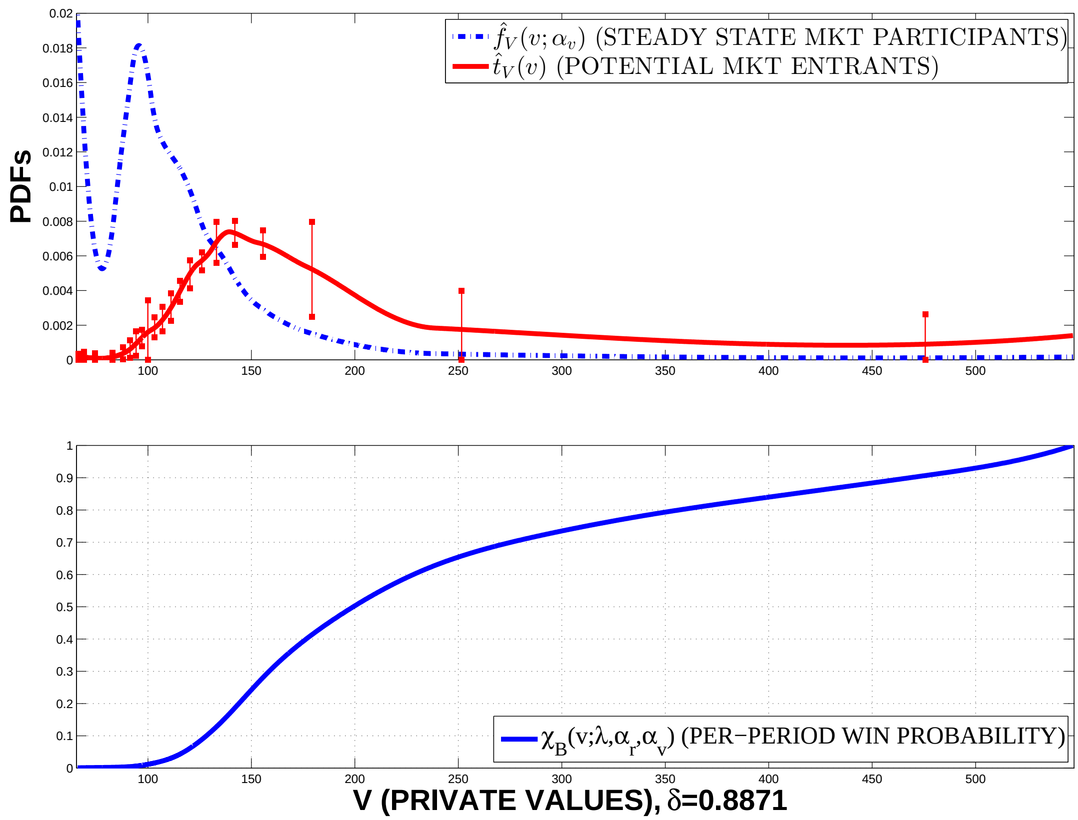

FIGURE 6

equilibrium bid distribution 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 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.

Table 3 displays various descriptive statistics derived from Stage II estimates, including average private values, average private values of winners, and information rents (i.e., the difference between the winner's private value and the spot-market price). The last row of the table shows information rents as a fraction of the winner's private value, on average. Finally, Figure 6 presents other Stage II estimates related to the distribution of buyer values. The upper pane displays the PDF of the distribution of market participants' private values in steady state under our preferred specification, $\hat{f}_V(v;\boldsymbol{\alpha}_v,\delta = 0.8871)$ (dash-dot line), as well as 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 v for comparison. Although there are many buyers in the market with low values in steady state, our model suggests that relatively few of these agents enter the market each period. However, those low-value buyers that do enter must stay in the market for a long period of time before winning, as indicated by the function χ. On a related note, the delay between entry and trade implies that the lower-value buyers are the ones most significantly affected by the per-period participation cost, $\hat{\kappa} = \$0.065$.

In comparing estimates of f_V and t_V in Figure 6, two differences should be noted. First, since f_V depicts the type distribution for all market participants—including players remaining from previous periods after failing to win a spot-market auction—it represents a measure $\lambda_1/(1-\lambda_2) = 7.96$ of agents (recall that sellers are assumed to have measure 1). On the other hand, t_V describes the type distribution of the measure $\mu = 0.9649$ of new agents that enter the market each period in order to maintain the steady state.

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

5. COUNTERFACTUALS

We now perform two counterfactual analyses to investigate the economic implications of our structural model. The first explores market efficiency. The second decomposes the relative importance of what we refer to as *platform composition* (PC) effects (i.e., market entry/exit when market conditions change) and *dynamic incentive* (DI) effects (i.e., when bidding behavior changes in response to shifts in opportunity costs).

Before proceeding, we would like to briefly describe the algorithm used to compute the counterfactuals. For expositional clarity, we focus on the SPA pricing rule. The structural primitives of our model are μ, t_V, κ, G_R, λ_2, and the space of types $[\underline{v}, \overline{v}]$.[30] These structural parameters remain fixed when performing our counterfactual exercises unless otherwise noted. We need to pin down the endogenous variables e, β, C, λ_1, and F_V. Recall that in our dataset $e = \underline{v}$, but in the counterfactuals we consider we often find $e > \underline{v}$. Any

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

equilibrium of a market with a SPA spot market must satisfy the following conditions:

$$\mu t_V(v) = \chi(\beta(v)) f_V(v)\mathcal{C} \tag{35}$$

$$\beta(v) = \beta^s(v - \delta\mathcal{V}(v)) = v - \delta\mathcal{V}(v) \tag{36}$$

$$\mathcal{V}(v) = \sum_{t=0}^{\infty} \delta^t \int_e^v (1 - \chi(\beta(s)))^t \chi(\beta(s)) ds \tag{37}$$

$$\chi(\beta(e))e - \rho(\beta(e)) = \kappa \tag{38}$$

$$F_V(\overline{v}) = 1 \tag{39}$$

Equation 35 requires that the distribution and measure of buyers exiting after winning an auction equals the distribution and measure of buyers flowing in, which is the stationarity condition of Definition 2.4. Equation 36 pins down the bidding strategy, while equation 37 uses the envelope theorem to describe the value function. Equation 36 would use an ordinary differential equation in non-SPA spot markets to describe β^s, but this is a relatively trivial modification to our structure. Equation 38 requires that the lowest value buyer that enters be indifferent to entering, which is the optimal entry condition of Definition 2.4. Finally, Equation 39 requires that the steady-state distribution of types be properly normalized.

In practice, our software uses a bisection algorithm to search for the equilibrium value of e. Given e, we let the auction size parameter, λ_1, adjust so that Equation 38 holds. Since $\mathcal{C} = E[K] = \frac{\lambda_1}{1-\lambda_2}$ by assumption, we immediately have $\mathcal{C}$. Given e, $\boldsymbol{\lambda}$, and $\mathcal{C}$ we can solve Equations 35 - 37 to obtain candidate values for f_V, β, and $\mathcal{V}$. If equation 39 fails to hold for the candidate f_V (i.e., if the steady-state distribution of types is not properly normalized), then we adjust our guess for e and repeat the process.[31]

Remark 1 (Uniqueness of the Stationary Competitive Equilibrium)**.** *It is difficult to provide general sufficient conditions for a unique SCE in our setting (even for an SPA spot market) because analytically characterizing the influence of e on F_V is hard due to the presence of participation costs. However, equations (35)–(39) provide a simple numerical test for uniqueness. Since the only unknown endogenous variable is e, let $F_V^{\tilde{e}}(\overline{v}) \equiv \int_{\tilde{e}}^{\overline{v}} f_V(v)dv$ be defined as the un-normalized integral over bidder types that is attained by solving equations (35)–(38) for some candidate cutoff, $\tilde{e}$, without imposing the normalization in (39). Then one can easily compute this quantity for a grid of $\tilde{e}$ ranging over the baseline type support $[\underline{v}, \overline{v}]$ to see whether $F_V^{\tilde{e}}(\overline{v})$ is strictly monotone in $\tilde{e}$. When this condition is satisfied–and it is for the counterfactuals we compute–there is a unique value of e*

[31]Since we do not have any data on the distribution of buyer values below the $\underline{v}$ observed in the data, we can only perform counterfactuals that yield a value of $e \geq \underline{v}$.

that satisfies equations (35)–(39), which in turn means the equilibrium is unique.[32] *Fortunately, this numerical test is easy to implement and fast to compute, so verifying uniqueness is quite tractable.*

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

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

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

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

[32] In the case where $\kappa = 0$ as in Backus and Lewis [2016], $e = 0$ in equilibrium and equations (35)–(39) pin down the remaining endogenous variables, which implies the equilibrium of our model is unique.

5.1.2. *Structural Welfare Calculations.* Our full Stage II structural estimates allow us to get a more complete idea of the frequency and magnitude of market inefficiency. First, using equation 6 we can compute the precise frequency of inefficient allocations as the fraction of all transactions involving low-value bidders:

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

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

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

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

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

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

5.1.3. *Counterfactual Market Centralization.* We now consider the extent to which inefficiencies can be mitigated by changing the market structure to one in which the same number of Kindles are allocated each period, but using fewer u-unit, uniform-price auctions with $u \geq 2$. Since new Kindles are relatively homogenous products, we think it is reasonable to assume that buyers view them as nearly perfect substitutes for one another. This suggests that our proposal to take steps toward more efficient market centralization using multi-unit

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

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

auctions is feasible. For products that 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.

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

$$\mathrm{E}[K_u] = \frac{\lambda_{1,u}}{1-\lambda_2} = u\mathcal{C} \tag{40}$$

Equation 40 mandates that the (endogenous) measure $\mathcal{C}$ of bidders be exactly assigned to the $1/u$ measure of u-unit auctions, much as in part 1 of Assumption 2.1. We assume λ_2, the dispersion parameter, is fixed at the estimated value and allow $\lambda_{1,u}$, the size parameter, to adjust so equation (40) is satisfied in our counterfactuals.

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

$$\chi_u(b) = G_R(b)\left[\sum_{m=0}^{u-1}\pi_M(m) + \sum_{m=u}^{\infty}\pi_M(m)\sum_{i=0}^{u-2}\binom{m}{m-i}F_V(v)^{m-i}(1-F_V(v)^i\right]. \tag{41}$$

which we use with equations (35)–(40) to compute the SCE for a u-unit auction.

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

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

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

One might naturally expect that if eBay could re-design their platform market to increase allocative efficiency, then it ought to be able to benefit by capturing some of the increased gains from trade.[33] From a static perspective of a single auction, the revenue equivalence theorem implies that the choice of payment rule has no effect on the revenues of the seller. If the revenues change with the altered auction format it must be due to a combination of (1) the new allocation rule and (2) the change in the measure and value distribution of the bidders. A careful examination of the moving parts within the model indicates that the sign of these effects on revenue is ambiguous. On the one hand, a given bidder with a value above the new participation cutoff e_u faces fewer competitors in the market. On the other hand, her remaining competitors also value the object more highly on average. This complex combination of effects make it difficult to derive a priori predictions on bidding behavior and the resulting effects of revenue. Table 5, which describes the mean revenues generated per-auction as a function of u, demonstrates that the average sale price actually falls as u increases.

To help explain why revenue drops as efficiency rises, Figure 7 plots the probability of winning for each type of agent and the equilibrium bid function for the $u = 1$ (solid line) and $u = 8$ (dashed line) market structures. The win probability plot reveals that for most agents (especially those most likely to win), increasing market efficiency raises the probability that they will win a spot-market auction within a given period. This raises their

[33]Of course, there are some other examples where this general intuition does not hold; for example, the literature on optimal auctions suggests that efficiency-reducing reservation prices can increase revenue. It is not clear if/how this result applies to our counterfactual other than the general sense that allocative efficiency and revenue are in tension.

TABLE 5. Counterfactual Mean Auction Revenues

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

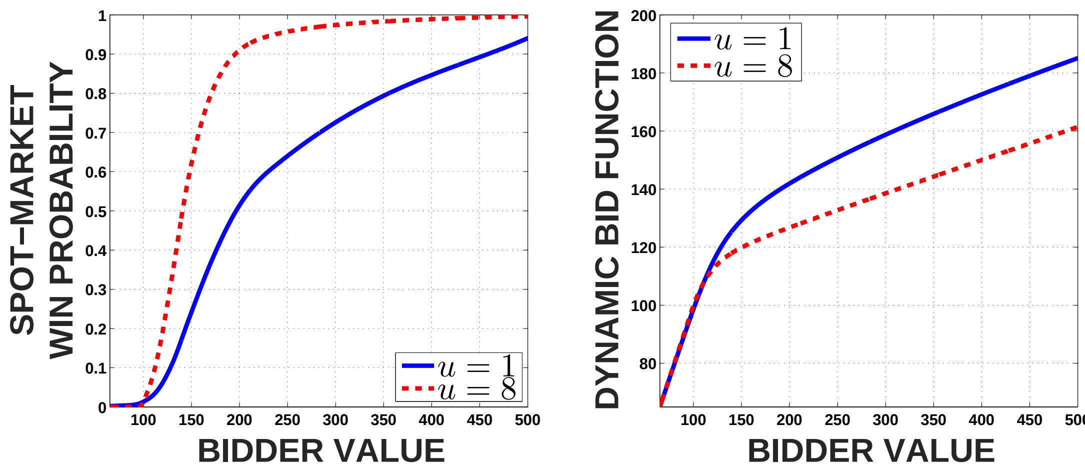

FIGURE 7. The Efficiency-Revenue Link

continuation values and therefore the opportunity cost of winning an auction today. This in turn reduces their bids by promoting further demand shading as shown in the second panel of Figure 7. Reduced bids then translate into decreased revenues for both sellers and eBay, which currently charges the sellers a percentage commission 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.

In Appendix D.2 we analyze the effect of centralization on the participation costs paid by the agents. Centralizing from single-unit auctions to 8-unit auctions reduces the participation costs paid by roughly 60%. Using a daily discount factor of $\delta = 0.95$, this amounts to a savings of $6.40 per Kindle, per day. There are two channels for the reduction in participation costs. First, because the good is allocated more efficiently when $u = 8$, high-valuation buyers win the good more quickly. Second, and more importantly, since low-valuation buyers are less likely to win the good each period, these agents leave the market and do not incur participation costs. Inducing these buyer-types to exit has a particularly strong effect on participation costs since low-valuation buyers stay in the market for many periods before winning when $u = 1$.

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

There are two effects when participation costs increase. First, agents' continuation values drop, which in turn reduces demand shading and increases bids. Holding the starting price distribution G_R fixed, these *dynamic incentive* (DI) effects increase allocative efficiency since bids are now more likely to exceed the starting price R. Second, an increase of the participation cost drives low-value buyers out of the market, which reduces the buyer-seller ratio and strengthens the steady-state distribution of active bidder types. The consequences of buyer selection out of the market are referred to as *platform composition* (PC) effects.

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

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

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

If all of the κ subscripts take on the same value, then χ_κ is generated by the SCE for that particular value of κ.

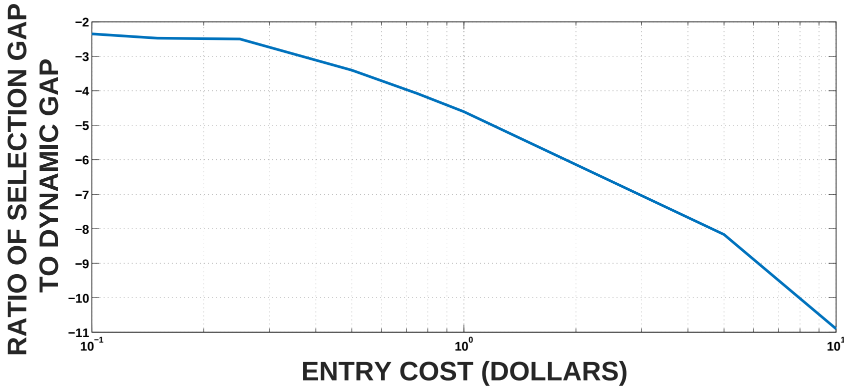

FIGURE 8. Relative Size of Dynamic and Selection Gaps

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

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

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

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

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

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 value function ($\mathcal{V}_{\kappa'}$) and bidding strategy ($\beta_{\kappa'}$) from an SCE with a higher participation cost κ'. We hold fixed the endogenous quantities that correspond to PC effects ($\boldsymbol{\lambda}$ and F_V) while allowing DI effects (β and $\mathcal{V}$) to vary with κ.

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

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

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

In Figure 8 we plot the ratio of the platform gap to the dynamic gap. When participation costs are low, the platform gap is twice as large as the dynamic gap. However, as costs rise, the platform gap becomes as much as ten times larger than the dynamic gap. The primary take-away from Figure 8 is that understanding the platform composition effects of market changes is more important than understanding the dynamic incentive effects of the changes. In terms of efficiency, the incentives driving selection into the platform market are at least as important as the incentives driving behavior once a buyer has chosen to participate.

6. CONCLUSION

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

Most platform markets exist in order to eliminate barriers to trade and allow for buyers and sellers to interact in a relatively low-friction environment. However, the sheer size of the markets may give rise to search frictions which prevent market outcomes from attaining the social optimum. We have estimated our model within the context of the market for Kindle Fire tablets, and we use these estimates both to compute the welfare loss under the present design and to suggest novel designs to mitigate these welfare losses. We begin by providing a "model-anemic" analysis that relies only on the bid distribution. We find a lower bound of at least 28% of auctions that close with a highest losing bidder whose private value exceeds that of some winner from another auction on that same day.

We then use our structural estimates to put a precise value on the deadweight loss and to study alternative spot-market mechanisms that eliminate (some of) the welfare loss due to search frictions. We find that over 36% of the auctions within a day allocate goods to winners with inefficiently low private values.[34] This results in a 14% welfare loss that can be attributed to the decentralized nature of the platform. This outcome implies that the single-unit auction market attains three quarters of total possible welfare improvement over a pure lottery system. By taking small steps toward a more centralized market structure—such as running multi-unit, uniform-price auctions with as few as 4 units each—2/3 of the remaining welfare loss can be recovered.

Another conclusion of our analysis relates to the importance of intertemporal incentives. In online auction markets, bid shading is driven by the opportunity cost of winning today,

[34]As in our model anemic analysis, we mean that such a lower value buyer won the good while a higher value buyer lost some other auction that day.

which depends on three factors: market tightness (ratio of buyers to sellers), market composition (ratio of high-value buyers to low-value buyers), and time preferences. The dynamic incentive to shade one's bid is much larger in magnitude than the more commonly studied static bid shading incentives generated by nontruthful pricing mechanisms. In other words, understanding opportunity costs is more important for buyers than understanding how to strategically respond to nontruthful pricing mechanisms.

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

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

References

Susan Athey and Philip A. Haile. Identification of standard auction models. *Econometrica*, 70 (6):2107–2140, November 2002.

Ned Augenblick, Muriel Niederle, and Charles Sprenger. Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics*, forthcoming, 2016.

Eduardo M. Azevedo and Jacob D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, forthcoming, 2016.

Matthew Backus and Gregory Lewis. Dynamic demand estimation in auction markets. *NBER Working Paper*, No. 22375, 2016.

Patrick Bajari and Ali Hortaçsu. The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34:329–55, 2003.

Eric Budish. Sequencing and information revelation in auctions for imperfect substitutes: Understanding ebay's market design. *typescript, Harvard University*, 2008.

Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *American Economic Review*, 86 (1):180–194, March 1996.

Stephen Burks, Jeffrey Carpenter, Lorenz Götte, and Aldo Rustichini. Which measures of time preference best predict outcomes: Evidence from a large-scale field experiment. *Journal of Economic Behavior & Organization*, 84:308–320, 2012.

Xiaohong Chen. Large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics*, 6(76):5549–5632, 2007.

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.

S. Donald, H. Paarsch, and J. Robert. An empirical model of the multi-unit, sequencial, clock auction. *Journal of Applied Econometrics*, 21:1221–1247, 2006.

A. Edlin and C. Shannon. Strict monotonicity in comparative statics. *Journal of Economic Theory*, 81:201–219, 1998.

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.

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.

Jianhua Z. Huang. Local asymptotics for polynomial spline regression. *Annals of Statistics*, 31(5):1600–1635, 2003.

Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *Management Science*, 60(12):2949–2970, 2014.

Matthew O. Jackson and Alejandro M. Manelli. Approximately competitive equilibria in large finite economies. *Journal of Economic Theory*, 77(2):354–376, 1997.

M. Jofre-Bonet and M. Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):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.

P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica*, 64:157–181, 1994.

Roger B. Myerson. Population uncertainty and poisson games. *International Journal of Game Theory*, 27:375–392, 1998.

Denis Nekipelov. Entry deterrence and learning prevention on ebay. *mimeo*, 2007.

Yoshihiko Otani and Joseph Sicilian. Limit properties of equilibrium allocations of walrasian strategic games. *Journal of Economic Theory*, 51(2):295–312, 1990.

Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 2:601 – 653, 2014.

Michael Peters and Sergei Severinov. Internet auctions with many traders. *Journal of Economic Theory*, 130:220–245, 2006.

V. Petrov. *Limit Theorems of Probability Theory: Sequences of Independent Random Variables*. Oxford: Oxford University Press, 1995.

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.

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

Appendix A. ONLINE SUPPLEMENTAL MATERIALS FOR *How Efficient are Decentralized Auction Platforms?* by Aaron L. Bodoh-Creed, Jörn Boehnke, and Brent R. Hickman THEORY: TECHNICAL PROOFS APPENDIX

We break our proof appendix into parts that correspond to the different sections of the main text. Subsection A.1 presents some background results on the topologies we employ. Subsection A.2 proves our claim that an equilibrium of our model exists. Subsection A.3 provides the proofs of our identification propositions.

A.1. **Preliminary Results.** Before beginning our main arguments, we define the Kolmogorov metric for random variables over $\mathbb{R}$, denoted d_K, where F and G are the CDFs of the random variables:

$$d_K(F,G) = \sup_x \|F(x) - G(x)\|$$

When either G or F is continuous (i.e., the underlying measure is atomless), then d_K metrizes the weak-$*$ topology (Petrov [1995], p. 43).

Remark 2. *Consider a random variable $X : \Omega \to \mathbb{R}$ with CDF $F(x)$. For N i.i.d. realizations, $\{x_1, ..., x_N\}$, drawn from F, denote the N realization empirical CDF as $F_N(x)$. Then from the Glivenko-Cantelli theorem we have*

$$\sup_{x\in\mathbb{R}} \|F_N(x) - F(x)\| \to 0 \text{ almost surely as } N \to \infty$$

Unless noted otherwise, we refer to almost sure convergence under d_K when making statements about the convergence of measures or their respective CDFs.

A.2. **Proofs from Section 2.** The highlight of this section is an equilibrium existence result using a fixed point argument. Before offering our proof, we first prove Proposition 2.3, which provides several essential preliminary results.

Proposition 2.3. *The best-response strategy $(\widetilde{\beta}, \widetilde{\zeta})$ satisfies:*

1. *$\widetilde{\beta}(v) = \beta^s(v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma))$ is a best-response bidding strategy.*
2. *$\widetilde{\beta}(v)$ is increasing in v if $\chi(b)$ is increasing and uniquely defined for a set of v of Lebesgue measure 1.*[35]
3. *$\widetilde{\beta}(v)$ if $\chi(b)$ is increasing.*
4. *There exists a cutoff $\widetilde{e}$ such that $\widetilde{\zeta}(v) = 1$ if and only if $v \geq \widetilde{e}$.*

Proof. Part (1) summarizes the discussion of the relationship between the static and dynamic auctions given in Section 2.2.

[35] $\chi(b)$ is strictly increasing in our application since G_R has full support.

To prove part (2), note that if a buyer enters the market, her optimal bid is defined by:

$$\underset{b}{argmax}\ \chi[b](v - \delta\mathcal{V}(v, C, F_V, G_R|\sigma)) - \rho[b] \tag{42}$$

Lemma A.5 proves that $\frac{\partial\mathcal{V}(v,C,G_B,G_R|\tilde{\sigma})}{\partial v} \in [0,1)$ and strictly positive for types with a best-response choice of *Enter*. Since the probability of being allocated the good is increasing in b, equation 3 is supermodular in (v, b). This implies that the best response function, $\widetilde{\beta}(v)$, is increasing in v (Milgrom and Shannon [1994]). In this case $\widetilde{\beta}(v)$ can admit only a countable set of discontinuities, which represent points where the bidder is indifferent between two bids. Since the set of such v is countable, the set has Lebesgue measure 0. Therefore the set of v where the best-response is uniquely defined is of Lebesgue measure 1. To prove part (3), note that if $\chi(b)$ strictly increasing in b, then $\widetilde{\beta}(v)$ is strictly supermodular and hence strictly increasing (Edlin and Shannon [1998]).

To see that ζ has the claimed cutoff form, suppose that it is optimal for a buyer with value v to enter and bid b. Since such a buyer could receive 0 by staying out of the market, it must be that $\mathcal{V}(v, C, G_B, G_R|\tilde{\sigma}) \geq 0$. Since $\frac{\partial\mathcal{V}(v,C,G_B,G_R|\tilde{\sigma})}{\partial v} \in [0,1)$, it must be that a buyer of type $v' > v$ also finds it optimal to enter. Therefore ζ has the claimed cutoff form of part (4). □

Our equilibrium existence argument consists of several parts. First we argue that the state variables, (C, F_V, G_R), live in a compact space that we denote Γ. We must show that the best responses to the state variables and the transition operator for the state variables are continuous with respect to Γ. Recall that we impose the Kolmogorov topology on the space $\Delta([0,1])$, which is metrized by d_K (described above). When considering continuity with respect to multiple variables, we impose the product topology on the Cartesian product of the relevant spaces.

As part of our argument we must prove that the bid strategies we consider have, stated informally, both a lower and an upper bound on their derivatives so that the induced distributions of bids do not admit atoms. Stated formally, we need to show that the best response function is continuous mapping from the space $\mathbb{C}_M[0,1|\varphi]$, $\varphi \in (0,1)$, into itself, where $\mathbb{C}_M[0,1|\varphi]$ contains all continuous, strictly increasing mappings from $[0,1]$ into $[0,1]$ that satisfy:

$$f(v) - f(v') \in \left[\varphi(v - v'), \frac{v - v'}{\varphi}\right],\ \ \forall v > v'.$$

Since the functions in $\mathbb{C}_M[0,1|\varphi]$ are Lipschitz continuous with modulus of continuity $1/\varphi$, this set of functions is equicontinuous and bounded over a compact domain. From the Arzelá-Ascoli theorem, $\mathbb{C}_M[0,1|\varphi]$ is compact when $\mathbb{C}_M[0,1|\varphi]$ is endowed with the sup-norm topology. The strategy space we need to consider can be defined as $\sigma = (e, \beta) \in \Xi =$

$[0,1] \times \mathbb{C}_M[0,1|\varphi]$. Once we show that our best response and state transition operators are continuous, a straightforward fixed point argument applies.

Our proof characterizes two functions. First, we define the best response dynamics, $BR(C, F_V, G_R, \sigma) = \tilde{\sigma} = \left(\tilde{e}, \tilde{\beta}\right)$, that describe how each type of buyer responds given the belief that (C, F_V, G_R) is stationary and all other agents use the strategy σ. We show that BR is continuous in (C, F_V, G_R, σ) and maps into Ξ.

The function $T(C, F_V, G_R|\tilde{\sigma}) = (\tilde{C}, \tilde{F}_V, \tilde{G}_r)$ describes the transitions of the aggregate states. We show that T is continuous in $(C, F_V, G_R, \tilde{\sigma})$ for $(C, F_V, G_R) \in \Gamma$ and $\tilde{\sigma} \in \Xi$, and $T(\cdot|\tilde{\sigma}) : \Gamma \to \Gamma$ as long as $\tilde{\beta} \in \mathbb{C}_M[0,1|\varphi]$. Once we have proven T and BR are continuous operators over compact spaces (and hence have compact images), we define the total operator $\mathcal{L} : \Gamma \times \Xi \to \Gamma \times \Xi$ where

$$\mathcal{L}(C, F_V, G_R, \sigma) = (\tilde{C}, \tilde{F}_V, \tilde{G}_R, \tilde{\sigma})$$
$$BR(\sigma, C, F_V, G_R, \sigma) = \tilde{\sigma}$$
$$T(C, F_V, G_R|\tilde{\sigma}) = (\tilde{C}, \tilde{F}_V, \tilde{G}_R)$$

Since G_R is constant across periods, $G_R = \tilde{G}_R$. A straightforward application of Schauder's fixed point theorem to the operator $\mathcal{L}$ gives us existence of a stationary equilibrium. However, proving T and BR are continuous operators over compact spaces requires many small steps.

Several of our results require that F_V and G_R admit bounded PDFs, f_V and g_r, except at $r = 0$. For notational purposes, we let $\mathcal{Q}[0, \bar{q}]$ denote the space of measures over $[0,1]$ that admit pdfs bounded from above by $\bar{q} > 0$.

A.2.1. *Continuity of Best Responses.* In this section we prove that the best responses are continuous as required and lie in $\mathbb{C}_M[0,1|\varphi]$ for some choice of $\varphi \in (0,1)$. Our first result is that the measure of entering agents and the distribution of bids are continuous in the underlying economy. To prove this, we need the following intermediate result. We add the "open neighborhood" qualifier to the statement of Lemma A.1 so that our result applies to empirical measures that are close to a nonatomic measure with a bounded PDF. We exclusively work with measures with bounded PDFs in Section A.2, but Section B.3 requires us to work with empirical measures near the SCE steady-state distributions.

Lemma A.1. *Consider any increasing function $f \in \mathbb{C}_M[0,1|\varphi]$, $\varphi \in (0,1)$, and let Z be an atomless CDF over $\mathbb{R}$ that admits a pdf that is bounded from above by M. Then $Y(s) = Z(f^{-1}(s))$ is uniformly continuous in $f \in \mathbb{C}_M[0,1|\varphi]$, $\varphi \in (0,1)$ and an open neighborhood of Z.*

Proof. Consider $f, \tilde{f} \in \mathbb{C}_M[0,1|\varphi]$ where $\left\|f - \tilde{f}\right\| \leq \gamma$. We can then write:

$$f(x + \gamma/\varphi) > \tilde{f}(x) > f(x - \gamma/\varphi)$$

which implies that for any y we have $\widetilde{f}^{-1}(y) \in [f^{-1}(y) - \gamma/\varphi, f^{-1}(y) + \gamma/\varphi]$, which in turn implies

$$Z\left(\widetilde{f}^{-1}(y)\right) \in \left[Z\left(f^{-1}(y) - \gamma/\varphi\right), Z\left(f^{-1}(y) + \gamma/\varphi\right)\right]$$

Since the PDF of Z is bounded by M, Z can increase by at most $M\gamma/\varphi$ over the intervals $[f^{-1}(y) - \gamma/\varphi, f^{-1}(y)]$ or $[f^{-1}(y), f^{-1}(y) + \gamma/\varphi]$, and we then have:

$$\left\|Z\left(\widetilde{f}^{-1}(y)\right) - Z\left(f^{-1}(y)\right)\right\| < \frac{M\gamma}{\varphi}$$

Since this bound holds uniformly over x, our result regarding continuity with respect to f is proven. Continuity with respect to Z is immediate given our use of the metric d_K over the space of measures. □

Lemma A.1 yields the following result.

Lemma A.2. *The distribution of bids, G_B, and the measure of entering buyers, $\mathcal{C}$, are continuous in C, F_V, and $\sigma = (e, \beta)$ provided that the entry cutoff point $e < 1$, F_V admits a bounded pdf, and $\beta \in \mathbb{C}_M[0,1|\varphi], \varphi \in (0,1)$.*

Proof. The distribution of entering buyers is

$$F_V^E(x) = \frac{F_V(x) - F_V(e)}{1 - F_V(e)} \text{ for } x \geq e, 0 \text{ otherwise}$$

and the measure of entering buyers is $\mathcal{C} = C\,[1 - F_V(e)]$. Note that F_V^E and $\mathcal{C}$ are continuous in (C, F_V). The distributions of bids can be described as

$$G_B(b) = F_V^E(\beta^{-1}(b))$$

Using Lemma A.1, we find that G_B is continuous in β. Finally, if F_V is atomless, then $F_V(e)$ varies continuously in e, so $(\mathcal{C}, G_B)$ is also continuous in e. □

Since we can state many of our lemmas most concisely in terms of continuity with respect to χ, we now prove that χ is continuous with respect to the underlying aggregate states and the agent's bid.

Lemma A.3. *Assume G_B is atomless. Then χ is continuous in $\boldsymbol{\lambda}$, G_R, G_B, and b. χ is strictly increasing in b.*

Proof. A bidder's beliefs about the number of other agents assigned to her auction are

$$\pi_M(k;\boldsymbol{\lambda}) = \pi(k;\boldsymbol{\lambda})\frac{(k+1)}{\mathrm{E}[K]} \text{ with } \mathrm{E}[K] = \mathcal{C}$$

Since $\mathcal{C}$ is continuous in C, F_V, and σ, $\pi_M(k;\boldsymbol{\lambda})$ is continuous with respect to these variables as well. The probability of winning the good (from a bidder's perspective) is

$$\chi(b) = G_R(b)\sum_{k=0}^{\infty}\pi_M(k;\boldsymbol{\lambda})G_B(b)^k$$

Recall that we have imposed the Kolmogorov topology over the spaces of measures, which means that two measures are close if their respective CDFs are close in the sup-norm. Under this topology, $\chi(b)$ is continuous in $\boldsymbol{\lambda}$, G_R, and G_B. $\chi(\cdot)$ is clearly continuous in b if G_B is atomless. To see that χ is strictly increasing, it suffices to note that G_R has full support, which means χ is strictly increasing even if b is not in the support of G_B. □

Lemma A.3 would not be true if we had imposed the weak-* topology over the bid or starting price distributions. To see this, suppose that G_B had an atom of measure 1 at $b = 0.5$, G_B^* has an atom of measure 1 at $b = 0.5 + \varepsilon$, and $G_R(0) = 1$. Under the Levy-Prokhorov metrization of the weak-* topology, G_B and G_B^* differ by ε. However, a bid of $b = 0.5 + \varepsilon/2$ would win with certainty under G_B and lose with certainty under G_B^*. In other words, χ fails to be continuous with respect to G_B. In contrast, under the Kolmogorov metric G_B and G_B^* differ by 1, so continuity is not threatened.

We required the assumption that $e < 1$ in order to prove Lemma A.2. Our next lemma proves that we can restrict attention to $e \in [0, \bar{e}]$ where $\bar{e} < 1$. The basic idea is that for large e, any entering buyer effectively bids against the distribution of starting prices. If κ is small, then buyers with values close to $v = 1$ find it strictly optimal to enter the market. This in turn implies that the best response e is always strictly less than 1. For the duration, we assume that any entry cutoff, e, is drawn from the set $[0, \bar{e}]$ where $\bar{e} < 1$.

Lemma A.4. *Fix G_R. Then as $e \to 1$, a buyer with type $v = 1$ strictly prefers to enter for $\kappa > 0$ sufficiently small.*

Proof. As $e \to 1$, the measure of entering buyers approaches 0. Assumption 2.2 implies that as $\mathcal{C} \to 0$ we have $\pi_M(0;\lambda) \to 1$. Thus, the payoff for a bidder that chooses a bid of b is:

$$G_R(b)v - \rho(b) \geq G_R(b)(v - b)$$

Since G_R has full support over $[0,1]$, $G_R(b) > 0$ for all $b > 0$, which means $G_R(b)(v-b) > 0$ for $b < v$. Therefore, the best response for buyers with v sufficiently large is to enter and bid $b < v$ as $e \to 1$. □

This result on the continuity of the bids combined with Assumption 2.5 yields the following result.

Lemma A.5. *$BR\left(C, F_V, G_R, \sigma\right) = \widetilde{\sigma} = \left(\widetilde{e}, \widetilde{\beta}\right)$ is continuous in (C, F_V, G_R) and σ if F_V admits a bounded PDF. Moreover, $\widetilde{\beta} \in \mathbb{C}_M[0,1|\widetilde{\varphi}]$, $\widetilde{\varphi} \in (0,1)$.*

Proof. Let $\mathcal{V}(v_i, C, G_B, G_R|\tilde{\sigma})$ denote the value function generated by best responding to (C, G_B, G_R) given a value of v_i. Theorem 1 of Pavan, Segal, and Toikka [2014] implies that $\mathcal{V}$ is almost everywhere differentiable and the derivative, where it exists, takes the value:

$$\frac{\partial \mathcal{V}(v, C, G_B, G_R|\tilde{\sigma})}{\partial v} = \sum_{\tau=t}^{\infty} \delta^{\tau-t}(1-\chi(\beta(v)))^{\tau-t}\chi(\beta(v)) \text{ if } v \geq e \tag{43}$$

and $\partial \mathcal{V}(v, C, G_B, G_R|\tilde{\sigma})/\partial v = 0$ if $v < e$. Since the strategies are best responses, the probability of sale, $\chi(b)$, must be positive for all types that enter. Since equation 43 is the sum of the probability of disjoint events with the sum exponentially discounted by δ, equation 43 is strictly bounded between 0 and 1. From the continuity of χ, we know that $\mathcal{V}$ is continuous in $(C, G_B, G_R, \tilde{\sigma})$

Given the continuity of G_B and $\boldsymbol{\lambda}$ with respect to (C, F_V, G_R) and σ, Assumption 2.5 implies that the spot-market strategy that is a best response to (C, F_V, G_R), $\widetilde{\beta}^s(v)$, is continuous with respect to (C, F_V, G_R) and v. Since $\mathcal{V}$ is continuous in these variables, the bidding strategy in our dynamic market, $\widetilde{\beta}(v) = \widetilde{\beta}^s(v - \delta\mathcal{V}(v, C, G_B, G_R|\tilde{\sigma}))$, is continuous in these variables as well.

Recall from Assumption 2.5 that $\widetilde{\beta}^s \in \mathbb{C}_M[0,1|\widetilde{\varphi}], \widetilde{\varphi} \in (0,1)$, which implies that $\widetilde{\beta}^s$ is almost everywhere differentiable. Where the derivative of $\widetilde{\beta}^s$ exists we can write:

$$\frac{d\widetilde{\beta}(v)}{dv} = \frac{d\widetilde{\beta}^s(v)}{dv}\left(1 - \delta\frac{\partial \mathcal{V}(v, C, G_B, G_R|\tilde{\sigma})}{\partial v}\right) \in \left(\varphi(1-\delta), \frac{1}{\varphi}\right) \tag{44}$$

Therefore $\widetilde{\beta} \in \mathbb{C}_M[0,1|\widetilde{\varphi}], \widetilde{\varphi}$ where $\widetilde{\varphi} = \varphi(1-\delta)$.

Since we know that the lowest valuation buyer willing to enter (i.e., $v = e$) has a continuation value of 0, it must be the case that:

$$\mathcal{V}(\tilde{e}, C, G_B, G_R|\tilde{\sigma}) = 0 \tag{45}$$

Since $\mathcal{V}(\circ, C, G_B, G_R|\tilde{\sigma})$ is continuous in (C, F_V, G_R) and σ and strictly increasing in v, the cutoff $\widetilde{e}$ is continuous in (C, F_V, G_R) and σ. □

A.2.2. *Continuity of Aggregate States.* Now we turn to proving that the evolution of the aggregate states is continuous. Before doing so, let us describe the operator $T(C, F_V, G_R|\tilde{\sigma}) = (\tilde{C}, \tilde{F}_V, \tilde{G}_R)$ in more depth. We iterate the states by using steady-state relations that, in equilibrium, are consistent with the aggregate state. We will need to use the following function, which the reader is encouraged to think of as an unnormalized version of $\tilde{f}_V$

$$\tilde{j}(v) = \frac{\mu t_V(v)}{\chi\left[\tilde{\beta}(v)\right]} \text{ if } \tilde{e} \leq v$$

where $\widetilde{e}$ is the best-response entry threshold and $\tilde{j}(v) = 0$ if $\tilde{e} > v$. We can define $\tilde{C}$ as

$$\tilde{C} = \int_0^1 \tilde{j}(s)ds$$

The distributions of buyer types and starting prices are

$$\tilde{f}_V(v) = \frac{\tilde{j}(v)}{\tilde{C}}$$
$$\tilde{G}_R = G_R$$

There is, at this point, no assurance that $\tilde{f}_V$ is well-defined since χ could have arbitrarily low values for some b. We rule this problem out with the following lemma.

Lemma A.6. $\chi(b) \geq \kappa > 0$ *if the agent is best-responding.*

Proof. For a buyer to find it optimal to enter given a positive entry cost, his expected payoff from entering must at least cover the cost of entry. This implies:

$$\chi v - \rho \geq \kappa$$

Therefore

$$\chi = \Pr\{\text{Transaction}\} \geq \frac{\kappa}{v} \geq \kappa$$

where the last inequality follows from the fact that $v \in [0, 1]$. □

Lemma A.6 implies that

$$\tilde{j}(v) \leq \frac{\mu t_V(v)}{\kappa}$$

for entrants. These relations give us

$$\tilde{C} \in [\mu, \mu/\kappa]$$
$$\tilde{f}_V(v) \in \left[0, \frac{t_V(v)}{\kappa}\right]$$

Therefore we can restrict attention to $\tilde{f}_V \in \mathcal{Q}\,[0, \bar{q}]$ where $\bar{q}$ is the maximum of $t_V(\cdot)/\kappa$. In other words, $\tilde{f}_V$ admits a bounded PDF. Furthermore, $\tilde{f}_V$ is continuous in χ as required. Finally, $\tilde{C}$, $\tilde{F}_V$ and $\tilde{G}_R$ inherit continuity with respect to $\boldsymbol{\lambda}$, G_R, G_B, and σ from the continuity of χ.

A.2.3. *Existence Result.* We can now prove the existence of a stationary equilibrium of the continuum model.

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

Proof. Lemmas A.1 through A.6 prove that $\mathcal{L}$ is a continuous mapping from $\Xi \times \Gamma$ into itself. Given the continuity of the mapping and the compactness of the spaces, we know that $\mathcal{L}$ has a compact image. Schauder's fixed point theorem implies that there exists a fixed point that defines a stationary equilibrium of our model. Lemma A.4 establishes that $e < 1$ (i.e., some buyers enter th market) for $\kappa > 0$ sufficiently small. □

A.3. **Identificatoin Proofs of Section 3.1.** Now we present the proofs of Lemma 3.2 and Proposition 3.3, which prove that our model is identified.

Lemma 3.2. *Suppose that the market tightness parameters take the form $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\}$, where $\lambda_k \equiv \pi(k; \boldsymbol{\lambda})$ for each $k = 0, 1, 2, \ldots$ Suppose further that all but I of them are zero, so that we can re-express them as $\boldsymbol{\lambda} = \{\lambda_{k_1}, \lambda_{k_2}, \ldots, \lambda_{k_I}\}$, where $k_1 < k_2 < \cdots < k_I$. Then it follows that there is a unique $(\boldsymbol{\lambda}, G_B)$ pair that is consistent with the joint distribution of the observables $(\tilde{K}, R, Y)$ for any finite I.*

Proof. If we condition on the event $\tilde{k} = k_I$ (i.e., no participants were filtered out) we then have $k = \widetilde{k} = k_I$. Therefore, we can write the distribution of the highest-losing-bid conditional on this event and reservation price r as:

$$H(b|r, k = k_I) = k_I G_B(b|b \geq r)^{k_I - 1} \left[1 - G(b_B|b \geq r)\right] + G_B(b|b \geq r)^{k_I}.$$

Since the above relationship is bijective (see explanation in Section 3.1.2) and since $H(b|r, k = k_I)$ is known, we can invert it to obtain $G_B(\cdot|b \geq r)$ for any r. Since

$$G_B(b|b \geq r) = \frac{G_B(b) - G_B(r)}{1 - G_B(r)} \text{ for } b \geq r$$

we can recover the form of the unconditional bid distribution $G_B(b)$ above r. Moreover, since the theoretical model implies buyers only enter the market if their bid might win an auction (i.e., $b \geq 0$), the infimum of the support of r must be weakly lower than the infimum of the bid support, and therefore we can recover the parent distribution of bids $G_B(b)$ on its entire domain.[36]

Having identified G_B also uniquely pins down a vector of probabilities $\boldsymbol{\lambda}$. To see why, consider the following equation:

$$\Pr[\tilde{k} = 0|r] = \sum_{k=0}^{k=k_I} G_B(r)^k \lambda_k. \tag{46}$$

For any choice of $(r_1, ..., r_I)$ where $G_B(r_i) \neq G_B(r_j), i \neq j$, the resulting I equations are linearly independent, which means only a single configuration of $\boldsymbol{\lambda}$ can be consistent with G_B and the joint distribution of the observables. Since Assumption 2.5 ensures that β is strictly increasing and we have assumed F_V is continuous, then for any $r_j > r_i$ where $G_B(r_i) \in (0, 1)$ implies $G_B(r_i) < G_B(r_j)$. Therefore we can choose from an uncountable set of conditions of the form of equation 46 to pin down our finite-dimensional parameter vector $\boldsymbol{\lambda}$. Finally, since the logic of the proof does not depend on the value of I, then by induction it is true for all finite I.

□

[36]To see this, let $\underline{r}$ be the infimum of the support of G_R. Conditional on the event $r = \underline{r}$ we have $G(b|b \geq \underline{r}) = \frac{G_B(b) - 0}{1} = G_B(b)$.

Proposition 3.3. *For a given discount factor δ, suppose that the market tightness parameters take the form $\boldsymbol{\lambda} = \{\lambda_0, \lambda_1, \lambda_2, \ldots\}$, where $\lambda_k \equiv \pi(k; \boldsymbol{\lambda})$ for each $k = 0, 1, 2, \ldots$ Suppose further that all but I of them are zero, so that we can re-express them as $\boldsymbol{\lambda} = \{\lambda_{k_1}, \lambda_{k_2}, \ldots, \lambda_{k_I}\}$, where $k_1 < k_2 < \cdots < k_I$. Then when the spot-market mechanism is a sealed-bid, second-price auction, for any finite I there is a unique configuration of model parameters $\boldsymbol{\Theta} \equiv (\boldsymbol{\lambda}, \kappa, F_V, \mu, T_V)$ that is consistent with the joint distribution of the observables $(\tilde{K}, R, Y)$.*

Proof. First note that Lemma 3.2 establishes that for any finite I there is a unique $(\boldsymbol{\lambda}, G_B)$ pair that is consistent with the joint distribution of the observables, and $G_R(r)$ is known. Given these three pieces, it also follows that the win probability $\chi(b)$ and the expected winner payment $\rho(b)$ are identified through equations (9) and (10) above. To identify the participation cost, combine equation (11) with the zero surplus condition (3.1) to find the following relation:

$$\chi(\underline{v})\underline{v} - \rho(\underline{v}) = \kappa \tag{47}$$

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 (13) 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:

$$\begin{aligned} \mu &= \int_{\underline{v}}^{\overline{v}} \chi\left[\beta(v)\right] f_V(v) dv \\ &= \left[1 - \pi(0)\right] G_R(\underline{b}) + \int_{\underline{b}}^{\overline{b}} g_R(r) \left(\sum_{k=1}^{\infty} \pi(k) \left[1 - G_B(r)^k\right] \right) dr. \end{aligned} \tag{48}$$

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

Appendix B. ONLINE SUPPLEMENTAL MATERIALS FOR *How Efficient are Decentralized Auction Platforms?* by Aaron L. Bodoh-Creed, Jörn Boehnke, and Brent R. Hickman THEORY: APPROXIMATING A FINITE MODEL

We refer to a model with a finite number of buyers and sellers as a *finite model*. Since the real world is clearly finite, we ideally would have estimated and computed counterfactuals using a finite model. The goal of this section is to justify our use of the continuum model as an approximation of the more realistic finite setting. We first lay out the primitives of the finite model analog to the continuum model described in Section 2. We then prove that an SCE of the continuum model is an approximate equilibrium of the finite model. This approximation result is our justification for the use of a continuum model as a proxy for the intractable, but more realistic, finite model.

B.1. **Primitives of the Finite Model.** We consider a sequence of games indexed by N where N refers to the number of sellers that list goods for sale in each period. All variables pertaining to the N-agent game are superscripted with N.[37] Each seller has a starting price R that is drawn randomly from the distribution G_R, which yields an empirical distribution of starting prices equal to $G^N_{R,t}$ in period t. The numbers of potential entrant buyers at $t = 0$ is denoted C^N_0. We assume that as $N \to \infty$

$$\frac{C^N_0}{N} \to C \in \mathbb{R}_{++}.$$

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

We now describe the stochastic matching process that assigns bidders to auctions in the finite setting. We denote the number of buyers that enter the market in period t as $\mathcal{C}^N_t$. 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 $\pi(K; \boldsymbol{\lambda})$.

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. We refer to a seller as being *seller rationed* if the supply of buyers is not sufficient to provide that seller with the number of buyers they are allocated by the matching process. Conditional on being matched, a particular bidder wins her auction if

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

$$\frac{S^N_t}{N} \to 1 \text{ almost surely as } N \to \infty$$

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

her bid is larger than the maximum of all competing bids and the seller's starting 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 starting price, then we assume the bidder wins the item.

We require the following assumption on the distribution of auction sizes in the finite game:

Assumption B.1. *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:*

$$\sqrt{NVar[K]}Pr\{Z_N = k\} \to \psi\left(\frac{k - NE[K]}{\sqrt{NVar[K]}}\right) \text{ uniformly over } k \in \mathbb{Z} \tag{49}$$

We use this assumption to approximate the probability mass function of sums of realizations of K using the probability density function of the normal distribution. Local limit theorems apply to most distributions of interest to economists including the generalized Poisson distribution used in our estimator. For the interested reader, Theorem 3.1 and Lemma 3.1 of McDonald McDonald [2005] provide conditions for the application of a local limit theorem, and these conditions amount to assuming that:

$$\sup_k N \; |PR\{Z_N = k+1\} - PR\{Z_N = k\}| \tag{50}$$

has a finite limit as $N \to \infty$.

Bidding strategies can be written as functions $\mathcal{O} : [0,1] \times \mathbb{R}_+ \times \Delta([0,1]) \times \Delta([0,1]) \to [0,1]$ with a typical bid denoted $\mathcal{O}(v_i, C_t^N, F_{V,t}^N, G_{R,t}^N)$. The entry decision for participating buyers is a function of the form $\theta : [0,1] \times \mathbb{R}_+ \times \Delta([0,1]) \times \Delta([0,1]) \to \{Enter, Out\}$ with a typical realization $\theta(v_i, C_t^N, F_{V,t}^N, G_{R,t}^N)$. We let Σ denote the buyers' strategy space. We let the bid and entry decision functions condition on the realized aggregate state, $\left(C_t^N, F_{V,t}^N, G_{R,t}^N\right)$, to account for fluctuations in the aggregate state or nonstationarities in the equilibrium. We continue to assume, as we did in the main text, that an agent that chooses *Out* exits the game immediately.

It is relatively easy to extend the notation used in the main text to nonstationary equilibria as we have for the finite model, and the working paper version of this project did so. Such an extension would be useful if one wished to, for example, study the effect of a firm's announcement of a release date for a new version of a consumer electronics product. In general, it is difficult to prove existence results for nonstationary models, but there are some cases which can be handled by simple extensions of our results. For example, suppose a company announces that a new product will be introduced T periods in the future. It is relatively easy to characterize a stationary equilibrium that applies after the introduction and characterize the equilibrium in the T periods leading up to the introduction using backwards induction. Another extension would be to assume that product introductions occur with a fixed probability in each period and are not announced in advance. To handle this case, one could include a state variable in the model that indicates the generation of product that exists in the market in the current period. If one assumes that innovation in the product ends after a fixed number of versions are released, then one could characterize the final stationary-state after the innovations have been exhausted and then use backwards induction to solve for the entire equilibrium.[38]

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

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

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

All agents discount future payoffs using a per-period discount factor $\delta \in (0,1)$. The value function given a (symmetric) SCE $\sigma = (e, \beta)$ for an agent with value v that bids $b = \beta(v)$ is:

$$\begin{aligned} \mathcal{V}^N(v, C_t^N, F_{V,t}^N, G_{R,t}^N|\sigma)) = &\mathbb{1}\{v \geq e\} \left(\chi^N[b]\left(v - \delta E_t^N\left[\mathcal{V}^N(v, C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N|\sigma)\right]\right) - \rho[b] - \kappa\right) \\ &+\delta E_t^N\left[\mathcal{V}^N(v, C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N|\sigma)\right]\Big) \end{aligned}$$

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

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

Definition B.2. The strategy vector σ and the initial state $C_0^N \in \mathbb{R}_{++}$ and $F_{V,0}^N, G_{R,0}^N \in \Delta_N([0,1])$ is an ε-Bayes-Nash Equilibrium (ε-BNE) of the N-agent game if for all bidder values v we have

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

Given the dynamic nature of our game, a solution concept that incorporates some notion of perfection might be expected. Consider the two ways in which an ε-BNE can yield an $\varepsilon > 0$. First, it may be that the agent does not exactly optimize with respect to high probability events, which results in a small loss with high probability. Second, the strategy may not optimize with respect to very rare events. Failing to optimize with respect to rare

[38]One could alternatively allow innovation to continue indefinitely and simply ignore innovations after some period T. This would result in an approximate equilibrium since the predictions of agent activity today would not account for innovations after period T. The upshot is that the equilibrium would be tractable to compute.

events can be approximately optimal but severely violate perfection. A stationary strategy can be an ε-BNE even though perfection may not even be approximately satisfied at the histories of the finite game in which the market aggregates differ significantly from the stationary state.

B.2. **Approximating the Large Finite Model.** It is not difficult to see why the model becomes too computationally complex to solve precisely as $N \to \infty$. In the N-agent game, the bidder's strategy must condition on all possible values of C_t^N, $F_{V,t}^N$, and $G_{R,t}^N$, which means the complexity of the strategies grows exponentially with N. The bidding strategy in the continuum need only condition on the values of C_t, $F_{V,t}$, and $G_{R,t}$, which evolve deterministically in equilibrium.

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

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

Proposition B.3. *Consider a SCE (σ, C, F_V, G_R). 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*

$$\left\|C_0^N - C\right\| + \left\|F_{V,0}^N - F_V\right\| + \left\|G_{R,0}^N - G_R\right\| < \eta \tag{51}$$

Proposition B.3 may be seen as providing an approximation to the actual equilibrium being played within the data-generating process, but it admits an alternative interpretation as a behavioral strategy. If one assumes that agents are subject to small computational costs, then in large markets it may be that they follow SCE behavioral predictions in lieu of solving a complex optimization problem for a vanishing benefit. Finally, note that while our result requires that the aggregate state in $t = 0$ be close to the SCE state, if we assume that seller and bidder types are drawn from F_V and G_R with numbers close to $N\,C$ and N, then $\left(C_0^N, F_{V,0}^N, G_{R,0}^N\right) \to (C, F_V, G_R)$ almost surely as $N \to \infty$. In other words, Equation 51 above is very likely to hold in large markets, and becomes increasingly so as $N \to \infty$.

B.3. **Proofs from Section B.2.** Our approximation result requires two steps. First, we must show that the limit game has a utility structure that is close to the utility structure of a sufficiently large finite game. In particular, we must show that the correlation across auctions vanishes as $N \to \infty$. Second, we must show that these facts imply that with high probability there are no deviations that yield a significant improvement in the utility of any agent. We conduct each task in separate sections. Throughout the sections we focus on an SCE strategy of the limit game (σ, C, F_V, G_R).

B.3.1. *Convergence of Utility.* First note that if an agent chooses *Out*, then his utility is 0 regardless of the number of other agents. For the remainder of the section we will assume the agent in question chooses *Enter.* The utility of a bidder in the current period of the N-agent game given the bidder enters and bids b is

$$\chi^N(b)v_i - \rho^N(b) - \kappa$$

If we can show that

$$\begin{aligned} \chi^N(b) &\to \chi(b) \\ \rho^N(b) &\to \rho(b) \end{aligned} \tag{52}$$

uniformly over b when we hold b, C, F_V, and G_R fixed, then we will have shown that the utility function in the N agent game converges to the utility functions of the limit game.

We first show that the probability of buyers being unmatched vanishes as $\mathcal{C}^N$ increases.

Lemma B.4. Pr (*A particular buyer is unmatched*) $= O\left(\frac{1}{\sqrt{N}}\right) + \omega_N$ *where* $\omega_N \to 0$ *as* $N \to \infty$.

Proof. Let D_l denote the number of bidders matched to auction l. For any buyers to be unmatched, the total "demand" for bidders from sellers must fall short of the supply, $\mathcal{C}^N$, which means that i buyers are not matched if and only if

$$\sum_{l=1}^{\mathcal{S}^N} D_l = \mathcal{C}^N - i$$

Since any buyer is equally likely to be amongst the unmatched buyers, conditional on i buyers being unmatched, the probability that a particular buyer is unmatched is $i/\mathcal{C}^N$. The total probability a particular buyer is unmatched is

$$\Pr(\text{A particular buyer is unmatched}) = \sum_{i=1}^{\mathcal{C}^N} \frac{i}{\mathcal{C}^N} \Pr\left[\sum_{l=1}^{N} D_l = \mathcal{C}^N - i\right]$$

Using assumption 2.1, we can approximate the probability mass function of the sum of the D_l using a normal distribution PDF, and the error ω_N vanishes as $N \to \infty$.[39] This lets us write:

$$\begin{aligned}
\sum_{i=1}^{\mathcal{C}^N} \frac{i}{\mathcal{C}^N} \Pr\left[\sum_{l=1}^{N} D_l = \mathcal{C}^N - i\right] &= \frac{1}{\mathcal{C}^N} \sum_{i=1}^{\mathcal{C}^N} \frac{i}{\sqrt{NVar[K]}} \psi\left[\frac{-i}{\sqrt{NVar[K]}}\right] + \omega_N \\
&\leq \frac{1}{\mathcal{C}^N} \frac{\mathcal{C}^N}{\sqrt{NVar[K]}} \psi[0] + \omega_N \\
&= \frac{\psi[0]}{\sqrt{NVar[K]}} + \omega_N = O\left(\frac{1}{\sqrt{N}}\right) + \omega_N
\end{aligned}$$

where ψ is the standard normal PDF. □

We first show that the probability that a seller is rationed vanishes as N increases.

Lemma B.5. $\Pr$ (*A particular seller is rationed*) $= O\left(\frac{1}{\sqrt{N}}\right) + \omega_N$ *where* $\omega_N \to 0$ *as* $N \to \infty$.

Proof. Let D_l denote the number of bidders matched to auction l. For any sellers to be rationed, the total "demand" for bidders from sellers must exceed the supply, $\mathcal{C}^N$, which means that the probability that exactly r sellers are rationed can be written:

$$\sum_{i=0}^{i=\infty} Pr\left[\sum_{l=1}^{N-r} D_l = \mathcal{C}^N - i\right] Pr\left[D_{N-r+1} \geq i\right]$$

Since any seller is equally likely to be amongst the rationed sellers, the probability that a particular seller is rationed is:

$$\sum_{r=1}^{N} \frac{r}{N} \sum_{i=0}^{i=\infty} Pr\left[\sum_{l=1}^{N-r} D_l = \mathcal{C}^N - i\right] Pr\left[D_{N-r+1} \geq i\right] \tag{53}$$

[39] For any approximation error $\epsilon > 0$ (i.e., the error in Equation 49), there exists $d < \infty$ independent of N or r such that if $N - r > d$ then the approximation error is less than ϵ. Since ϵ was chosen arbitrarily, the approximation error for these terms vanishes as $N \to \infty$. The terms in equation 53 where $N - r \leq d$, which may have a nontrivial error when using the local limit approximation, are a vanishing fraction of the terms in the outer summation of equation 53. For large N these also are the lowest probability terms in this summation since they represent the largest degree of seller rationing. Therefore, the total probability (and hence the total contribution to the approximation error) of the terms where $N - r \leq d$ vanishes as $N \to \infty$. Combining the arguments for the case where $N - r > d$ and the case where $N - r \leq d$, we have that the total approximation error vanishes as $N \to \infty$.

Using assumption 2.1, we can approximate the probability mass function of the sum of the D_l using a normal distribution PDF, and the error $\omega_N \to 0$ as $N \to \infty$. This lets us write:

$$\sum_{r=1}^{N} \frac{r}{N} \sum_{i=0}^{i=\infty} Pr\left[\sum_{l=1}^{N-r} D_l = \mathcal{C}^N - i\right] Pr\left[K \geq i\right] =$$

$$\sum_{r=1}^{N} \sum_{i=0}^{i=\infty} \frac{r}{N} \frac{Pr\left[K \geq i\right]}{\sqrt{(N-r)Var[K]}} \psi\left(\frac{rE[K]-i}{\sqrt{(N-r)Var[K]}}\right) + \omega_N$$

Using a change of variables we have:

$$= \sum_{i=0}^{i=\infty} \sum_{\underline{x}(i)=1}^{\overline{x}(i)} \frac{i + x\sqrt{NVar[K]}}{NE[K]} \frac{Pr\left[K \geq i\right]}{E[K]} \psi\left(x\right) \frac{E[K]}{\sqrt{(N-r)Var[K]}} + \omega_N \tag{54}$$

where

$$\underline{x}(i) = \frac{rE[K]-i}{\sqrt{(N-r)Var[K]}}$$

$$\overline{x}(i) = \frac{(N-1)E[K]-i}{\sqrt{Var[K]}}$$

We can decompose equation 54 into two terms. First consider:

$$\sum_{i=0}^{i=\infty} \frac{iPr\left[K \geq i\right]}{NE[K]^2} \sum_{\underline{x}(i)=1}^{\overline{x}(i)} \psi\left(x\right) \frac{E[K]}{\sqrt{(N-r)Var[K]}} \leq \sum_{i=0}^{i=\infty} \frac{iPr\left[K \geq i\right]}{NE[K]^2} \int_{-\infty}^{\infty} \psi\left(x\right) dx$$

$$= \sum_{i=0}^{i=\infty} \frac{iPr\left[K \geq i\right]}{NE[K]^2}$$

$$= \frac{1}{N}$$

The second term is:

$$\sum_{i=0}^{i=\infty} \sum_{\underline{x}(i)=1}^{\overline{x}(i)} \frac{x\sqrt{NVar[K]}}{NE[K]} \frac{Pr\left[K \geq i\right]}{E[K]} \psi\left(x\right) \frac{E[K]}{\sqrt{(N-r)Var[K]}}$$

$$\leq \frac{\sqrt{Var[K]}}{\sqrt{N}E[K]} \sum_{i=0}^{i=\infty} \frac{Pr\left[K \geq i\right]}{E[K]} \int_{-\infty}^{\infty} \|x\| \, \psi\left(x\right) dx$$

$$= \frac{\sqrt{Var[K]}}{\sqrt{N}E[K]} \int_{-\infty}^{\infty} \|x\| \, \psi\left(x\right) dx$$

$$= O\left(\frac{1}{\sqrt{N}}\right)$$

Combining these terms we find:

$$\sum_{r=1}^{N} \frac{r}{N} \sum_{i=0}^{i=\infty} Pr\left[\sum_{l=1}^{N-r} D_l = \mathcal{C}^N - i\right] Pr\left[D_{N-r+1} \geq i\right] = O\left(\frac{1}{\sqrt{N}}\right) + \omega_N$$

as required. □

The asymptotics results required to prove equation 52 holds are complicated by the fact that in the finite game the buyers are sampled without replacement when assigned to auctions. Intuition suggests that, much as in other settings with sampling without replacement, the covariance between two auctions ought to vanish as the number of auctions grows. We prove that this is indeed the case in Lemma B.6.

The proof of Lemma B.6 proceeds in two steps. First, we use the Chebyshev inequality to prove that the expected outcome of bidding b in the N-agent game, $\chi^N(b)$ and $\rho^N(b)$, approaches the average outcome a bid of b would have generated across the N auctions. The second step is to show that the expected outcome of the N-agent game approaches the expected outcome of the game with a continuum of agents. Both steps require grappling with the correlation of outcomes across auctions, which is the only difference in the game mechanics between the model with a finite set of buyers and the model with a continuum of buyers.

We analyze the covariance between auctions caused by a rejection sampling algorithm for generating samples without replacement. In the rejection algorithm, a sample of buyers for each auction is generated with replacement. If the sample violates the conditions of sampling without replacement, then the sample is rejected. This process is repeated until a sample is not rejected. We prove that the probability of rejection vanishes as N increases, which in turn means that the covariance between auctions under the sampling without replacement regime of the finite model converges to the 0 covariance of the sampling with replacement regime of the continuum model as N increases.

There are three events that cause a sample of buyers for two auctions to be rejected. First, it could be that one of the sellers is rationed because the demand for buyers outstrips the supply. Lemma B.5 proves that this effect vanishes at a rate of $N^{-0.5}$. Second, it could be that a particular buyer is unmatched. Lemma B.4 proves that this effect vanishes at a rate of $N^{-0.5}$. The third event that causes a sample to be rejected is that a bidder appears twice in the sample. We show that the probability of this event also vanishes at the rate of $N^{-0.5}$.

Lemma B.6. *Supose that all buyers follow some SCE of the limit game, $\sigma = (e, \beta)$. For any $\varepsilon, \gamma > 0$ we can choose N^* such that for any $N > N^*$ and any (C^N, F_V^N, G_R) we have*

$$Pr\left[\frac{1}{\mathcal{C}^N}\left\|\sum_{i=1}^{\mathcal{C}^N}\left[x_i^N(b) - \chi^N(b)\right]\right\| > \varepsilon\right] \;<\; \gamma \tag{55}$$

$$Pr\left[\frac{1}{\mathcal{C}^N}\left\|\sum_{i=1}^{\mathcal{C}^N}\left[p_i^N(b) - \rho^N(b)\right]\right\| > \varepsilon\right] \;<\; \gamma$$

$$\sup_{b\in[0,1]}\left\|\chi^N(b) - \chi(b)\right\| = O\left(\frac{1}{\sqrt{N}}\right) \tag{56}$$

$$\sup_{b\in[0,1]}\left\|\rho^N(b) - \rho(b)\right\| = O\left(\frac{1}{\sqrt{N}}\right)$$

The choice of ε and ρ can be chosen uniformly over (C^N, F_V^N, G_R).

Proof. We provide a proof for Equations 55 and 56, but essentially identical arguments suffice for the other results. For notational cleanliness, we provide a proof of the probability of large positive deviations, but the analogous result for large negative deviations is essentially identical.

From Chebyshev's inequality we have

$$Pr\left[\frac{1}{N}\sum_{i=1}^{N}\left(x_i^N(b) - \chi^N(b)\right) > \varepsilon\right] \;\leq\; \frac{1}{\varepsilon^2}Var\left(\frac{1}{N}\sum_{i=1}^{N}\left(x_i^N(b) - \chi^N(b)\right)\right) \tag{57}$$

$$= \frac{1}{\varepsilon^2}\frac{1}{N^2}\Big[\sum_{i=1}^{N} Var\left(x_i^N(b) - \chi^N(b)\right) + \tag{58}$$

$$2\sum_{i=2}^{N}\sum_{j=1}^{N} cov\left(x_i^N(b) - \chi^N(b), x_j^N(b) - \chi^N(b)\right)\Big]$$

Since $x_i^N(b)$ is bounded, we know

$$\frac{1}{N^2}\sum_{i=1}^{N} Var\left(x_i^N(b) - \chi^N(b)\right) = O(N^{-1}) \tag{59}$$

We now bound the covariance term by assessing the covariance that would be generated by using rejection sampling to generate a set of buyers for two auctions. A rejection sampling algorithm draws samples of buyers for the auctions using sampling with replacement (SWR) from the pool of buyers in the finite game, but the sample generated is rejected if it either (1) the buyer is unmatched, (2) one of the sellers is rationed, or (3) the sample of buyers contains two "copies" of the same agent. This process is repeated until a sample is not rejected. Since the covariance between auctions is 0 under sampling with replacement, any covariance between auctions must be generated by events in which a sample is rejected. Since the price and probability of winning are bounded, we can bound the covariance due to these rejection events by bounding the probability that the sample is rejected. In the following argument we prove that the probability a sample is rejected vanishes, so the covariance terms vanish in the limit as $N \to \infty$.

Lemmas B.4 and B.5 provide a uniform upper bound on the probability of events (1) and (2). Let $\mathcal{E}_S$ denote the event the auctions share a bidder in a SWR regime, which allows us to write

$$cov\left(x_i^N(b) - \chi^N(b), x_j^N(b) - \chi^N(b)\right) \leq \Pr\left[\mathcal{E}_S\right] + O\left(\frac{1}{\sqrt{N}}\right) + \omega_N$$

where the second term captures the probability of a buyer being unmatched or a seller being rationed and $\omega_N \to 0$ as $N \to \infty$.

Fix two auctions with m and n bidders. The probability the auctions do not share a bidder is[40]

$$\left(1 - \frac{1}{\mathcal{C}^N}\right)\left(1 - \frac{2}{\mathcal{C}^N}\right) ... \left(1 - \frac{m+n-1}{\mathcal{C}^N}\right) > \left(1 - \frac{m+n}{\mathcal{C}^N}\right)^{m+n}$$

The probability the auctions share an agent is no more than

$$\sum_{\{m,n: m+n \leq \mathcal{C}_N\}} \pi(m; \boldsymbol{\lambda}) \pi(n; \boldsymbol{\lambda}) \left[1 - \left(1 - \frac{m+n}{\mathcal{C}^N}\right)^{m+n}\right]$$

A binomial expansion yields:

$$\left(1 - \frac{m+n}{\mathcal{C}^N}\right)^{m+n} = 1 - \frac{(m+n)^2}{\mathcal{C}^N} + o\left(\mathcal{C}_N^{-1}\right)$$

[40]If we think of drawing the $m+n$ buyers in order, the second buyer chosen cannot have the same identity as the first (first term), the the third buyer chosen must not have the same identity as either of the first two (second term), etc.

Therefore

$$\begin{aligned}
\sum_{\{m,n:m+n\leq \mathcal{C}^N\}} \pi(m;\boldsymbol{\lambda})\pi(n;\boldsymbol{\lambda}) \left[1-\left(1-\frac{m+n}{\mathcal{C}^N}\right)^{m+n}\right] \\
&= \sum_{\{m,n:m+n\leq \mathcal{C}^N\}} \pi(m;\boldsymbol{\lambda})\pi(n;\boldsymbol{\lambda})\frac{(m+n)^2}{\mathcal{C}^N} + o\left(\frac{1}{\mathcal{C}_N}\right) \\
&= \sum_{\{m,n:m+n\leq \mathcal{C}^N\}} \pi(m;\boldsymbol{\lambda})\pi(n;\boldsymbol{\lambda})\frac{m^2+2mn+n^2}{\mathcal{C}^N} + o\left(\frac{1}{\mathcal{C}^N}\right) \\
&\leq \frac{2E[K^2]+2E[K]^2}{\mathcal{C}^N} + o\left(\frac{1}{\mathcal{C}^N}\right) \\
&= O\left(\frac{1}{\mathcal{C}^N}\right)
\end{aligned}$$

Putting all of these results together, we have that

$$\Pr\left[\mathcal{E}_S\right] = O\left(\frac{1}{\mathcal{C}^N}\right) = O\left(\frac{1}{N}\right)$$

where the final equality follows from the fact that $\mathcal{C}^N = \Theta(N)$.

Pulling our argument together we have

$$\frac{1}{\varepsilon^2}\frac{1}{N^2}\sum_{i=2}^{N}\sum_{j=1}^{i-1} cov\left(x_i^N(b)-\chi^N(b), x_j^N(b)-\chi^N(b)\right) = O\left(\frac{1}{\sqrt{N}}\right) + \omega_N, \omega_N \to 0 \text{ as } N \to \infty$$

Referring back to equation 57, we then have

$$Pr\left[\frac{1}{N}\sum_{i=1}^{N}\left(x_i^N(b)-\chi^N(b)\right) > \varepsilon\right] = O\left(\frac{1}{\sqrt{N}}\right) + \omega_N, \omega_N \to 0 \text{ as } N \to \infty$$

The final step is proving that χ^N and χ are close to one another. Note that the difference between these two functions is generated by the fact that χ^N is generated by constructing the set of auctions with a sampling with replacement process, buyers can be unmatched and sellers can be rationed. Our argument regarding the correction required to account for these differences implies:

$$\left\|\chi^N(b)-\chi(b)\right\| = O\left(\frac{1}{\sqrt{N}}\right)$$

The uniformity with respect to b follows from standard arguments based on the separability of the reals and both functions being monotone in b. The uniformity with respect to (C^N, F_V^N, G_R) follows from the fact that the proof is completely independent of these variables. □

The following lemma implies that the within-period utility of the game with a finite number of agents converges to the within-period utility function of the continuum limit game.

Lemma B.7. *Supose that all agents follow some SCE of the limit game, (e, β). For any $\varepsilon > 0$ we can choose N^* such that for any $N > N^*$ and any (C^N, F_V^N, G_R^N) we have*

$$\sup_{b,v\in[0,1]} \left\| \chi^N(b)v_i - \rho^N(b) - (\chi(b)v_i - \rho(b)) \right\| < \varepsilon \tag{60}$$

Proof. The result follows from Lemma B.6. □

B.3.2. *Mean Field Lemma.* The goal of this subsection is to prove that the evolution of the finite game approaches the deterministic evolution of the limit game as $N \to \infty$. We start by showing the initial distribution of agent types in the finite and limit game converges as $N \to \infty$.

Lemma B.8. *The empirical distribution of types and starting prices in period 0 of the N-agent game converges to F_V and G_R with a convergence rate of $O(N^{-0.5})$.*

Proof. Follows from Remark 2. □

We use time indices for the variables in the next proposition to make the evolution of the aggregate variables clear. Let $Q_N(C_t^N, F_{V,t}^N, F_{C,t}^N|\sigma) = (C_{t+1}^N, F_{V,t+1}^N, F_{C,t+1}^N)$ denote the aggregate state iterator in the N-agent game, where $(C_{t+1}^N, F_{V,t+1}^N, F_{C,t+1}^N)$ is a random variable.

Lemma B.9. *Consider a stationary SCE strategy $\sigma = (e, \beta)$ and aggregate state (C, F_V, G_R). For any $\eta, \gamma > 0$, we can choose N^* such that for all $N > N^*$ and $\left(\widehat{C}_t^N, \widehat{F}_{V,t}^N, \widehat{G}_{R,t}^N\right)$ such that*

$$\left\|\widehat{C}_t^N - C\right\| + \left\|\widehat{F}_{V,t}^N - F_V\right\| + \left\|\widehat{G}_{R,t}^N - G_R\right\| < \frac{\eta}{2}$$

we have the following with probability at least $1 - \gamma$

$$\left\|\widehat{C}_{t+1}^N - C\right\| + \left\|\widehat{F}_{V,t+1}^N - F_V\right\| + \left\|\widehat{G}_{R,t+1}^N - G_R\right\| < \eta \tag{61}$$

where $Q_N\left(\widehat{C}_t^N, \widehat{F}_{V,t}^N, \widehat{G}_{R,t}^N|\sigma_{-i}, \sigma_i'\right) = \left(\widehat{C}_{t+1}^N, \widehat{F}_{V,t+1}^N, \widehat{G}_{R,t+1}^N\right)$.

Proof. As an initial note, when any single agent deviates from σ in the limit game, no change in the aggregate variables occurs. When a single agent deviates in the finite game, it causes a change of at most $(\mathcal{C}^N)^{-1}$. Since $(\mathcal{C}^N)^{-1} \to 0$ as $N \to \infty$ in any equilibrium of the finite game, these deviations do not affect the convergence arguments presented below.

First, we have $G_{R,t+1}^N \to G_R$ by Remark 2. What remains is to show that $\frac{C_{t+1}^N}{N} \to C_{t+1}$ and $F_{V,t+1}^N \to F_{V,t+1}$. Focusing on the new potential entrants, there are $\lceil N\mu \rceil$ agents added to the game with types $(\tilde{v}_1, ..., \tilde{v}_{\lceil N\mu \rceil})$ drawn from T_V. $\frac{\lceil N\mu \rceil}{N} \to \mu$ as $N \to \infty$, so Remark 2 implies

$$\frac{1}{\lceil N\mu \rceil} \sum_{i=1}^{\lceil N\mu \rceil} \mathbf{1}\left\{v \geq \tilde{v}_i \geq e\right\} \to T_V(v) \text{ uniformly over } v$$

This means that the only thing we need to show is that the number and type distribution of agents continuing onto the next period in the finite game converges to the analogous measure and distribution in the limit game as $N \to \infty$. We now show that the measure and distribution of buyers that exit each period converges in probability and uniformly over v to the analogous quantities in the continuum model as $N \to \infty$. Lemma B.6 implies:

$$\frac{1}{C_t^N} \sum_{i=1}^{C_t^N} \mathbf{1}\left\{v_i \geq e\right\}(1 - x_i(\beta(v)) \quad \to \quad \int_0^1 \mathbf{1}\left\{v_i \geq e\right\}(1 - \chi(\beta(v)) dF_{V,t}(s)$$

$$\frac{1}{C_t^N} \sum_{i=1}^{C_t^N} \mathbf{1}\left\{v \geq v_i \geq e\right\}(1 - x_i(\beta(v)) \quad \to \quad \int_0^1 \mathbf{1}\left\{v \geq v_i \geq e\right\}(1 - \chi(\beta(v)) dF_{V,t}(s)$$

The first equation refers to the measure of buyers that continue to the next period, while the second equation refers to the distribution of the types of these buyers. Bringing these results together, we have $\left(C_{t+1}^N, F_{V,t+1}^N, G_{R,t+1}^N\right) \to (C_{t+1}, F_{V,t+1}, G_R)$ in probability as $N \to \infty$, which is equivalent to our desired result. □

Iterating Lemma B.9 immediately gives us the following:

Corollary B.10. *Consider a stationary SCE strategy σ and aggregate state (C, F_V, G_R). For any $\Delta, \gamma > 0$, we can choose $\eta > 0$ and N^* such that for all $N > N^*$ and $\left(C_t^N, F_{V,t}^N, G_{R,t}^N\right)$ such that*

$$\left\|C_t^N - C\right\| + \left\|F_{V,t}^N - F_V\right\| + \left\|G_{R,t}^N - G_R\right\| < \eta$$

we have for all $t' \in \{t, .., t+\tau\}$ with probability at least $1 - \gamma$

$$\left\|C_{t'}^N - C\right\| + \left\|F_{V,t'}^N - F_V\right\| + \left\|G_{R,t'}^N - G_R\right\| < \Delta \tag{62}$$

where for $k \in \{0, .., \tau - 1\}$ we define $Q_N\left(C_{t+k}, F_{V,t+k}^N, G_{R,t+k}^N | \sigma_{-i}, \sigma_i'\right) = \left(C_{t+k+1}^N, F_{V,t+k+1}^N, G_{R,t+k+1}^N\right)$.

B.3.3. *No Profitable Deviations.* In this subsection, we finally prove our approximation result. We start by proving that the limit model has a continuous per-period utility function.

Lemma B.11. *$\chi(b)v - \rho(b) - \kappa$ is continuous with respect to (C, F_V, G_R) and β when F_V admits a PDF that is bounded from above*

Proof. $\chi(b)$ was proven to be continuous in Lemma A.3. We can write

$$\rho(b) = \pi_M(1;\boldsymbol{\lambda}) \int_0^b u \, G_R(du) + \sum_{k=2}^{\infty} \pi(k;\boldsymbol{\lambda}) \int_0^b \int_0^b \max\{u,t\} G_R(du) G_B^{k-1}(dt)$$

Lemma A.2 implies that G_B is continuous in (C, F_V, G_R) and (e, β). Since the integrands are continuous and G_B and G_R are continuous, the resulting integral is continuous.[41] □

These results, together with the convergence of the utility functions, yields the following result on the convergence of value functions.

Lemma B.12. *Consider a stationary SCE strategy σ and aggregate state (C, F_V, G_R). For any $\epsilon, \gamma > 0$ we can choose $\eta > 0$ and N^* such that for all $N > N^*$ and $\left(C_0^N, F_{V,0}^N, G_{R,0}^N\right)$ such that*

$$\|C_0 - C\| + \left\|F_{V,0}^N - F_V\right\| + \left\|G_{R,0}^N - G_R\right\| < \eta \tag{63}$$

we have with probability at least $1-\gamma$

$$\text{For all } v, \ \left\|\mathcal{V}^N\left(v, C_0^N, F_{V,0}^N, G_{R,0}^N|\sigma\right) - \mathcal{V}(v, C, F_V, G_R|\sigma)\right\| < \varepsilon$$

Proof. First note that if $v < e$, then we are done since the buyer never enters the market and receives the same payoff in either game. For the duration we assume that $v \geq e$.

Let $E_0^N[x_t]$, etc. refer to an agent's expectation in period 0 about an event that occurs in period t of the finite game. We can write the value functions as

$$\mathcal{V}^N\left(v, C_0^N, F_{V,0}^N, G_{R,0}^N|\sigma\right) = \sum_{t=0}^{\infty} \delta^t E_0^N\left[x_t\, v - p_t - \kappa | C_0^N, F_{V,0}^N, G_{R,0}^N, \sigma\right]$$

$$\mathcal{V}(v, C, F_V, G_R|\sigma) = \sum_{t=0}^{\infty} \delta^t \left(\chi v - \rho - \kappa\right)$$

Choose T such that $\delta^T < \frac{\varepsilon}{3}$, and note:

$$\sum_{t=T}^{\infty} \delta^t \left[(x_t v - p_t - \kappa)\right] < (1-\delta)\delta^T v < \frac{\varepsilon}{3}$$

From hereon, we consider only the first T periods.

Lemma B.7 implies that for any sample path of $\left\{\left(C_t^N, F_{V,t}^N, G_{R,t}^N\right)\right\}_{t=0}^{\infty}$ we can choose an N^* sufficiently large so that for all $t \in \{0, .., T\}$

$$\sup_{b,v\in[0,1]} \left\|E_t^N\left[x_t^N(b)v - p_t^N(b) | C_t^N, F_{V,t}^N, G_{R,t}^N\right] - (\chi(b)v - \rho(b))\right\| < \varepsilon \tag{64}$$

[41] We have implicitly used the fact that the topology generated by d_K is finer than the weak* topology.

where χ and ρ are conditioned on $(C_t^N, F_{V,t}^N, G_{R,t}^N)$. Lemma B.10 implies that for any $\Delta, \gamma > 0$ we can choose η sufficiently small and N sufficiently large such that

$$\left\| C_{t+\tau}^N - C \right\| + \left\| F_{V,t+\tau}^N - F_V \right\| + \left\| G_{R,t+\tau}^N - G_R \right\| \leq \Delta \tag{65}$$

for all $\tau \leq T$ with probability at least $1 - \gamma$.

From Lemma B.11, if Equation 65 holds and Δ is sufficiently small, we have for all $t \in \{0, .., T\}$ we have:

$$\left\| E_0 \left[xv - p - \kappa | C_t^N, F_{V,t}^N, G_{R,t}^N, \sigma \right] - (\chi(b)v - \rho(b)) \right\| < \frac{\varepsilon}{3T}$$

where χ and ρ are conditioned on the steady-state aggregate variables, (C, F_V, G_R). Note that this result holds uniformly over v and the closed neighborhood defined by Equation 65. Finally, in the complementary event that the sample path of $\left\{ \left(C_t^N, F_{V,t}^N, G_{R,t}^N \right) \right\}_{t=0}^{\infty}$ is not close to (C, F_V, G_R) (i.e., Equation 65 fails to hold for Δ sufficiently small) we have:

$$\left\| E_0^N \left[(xv - p - \kappa) | C_t^N, F_{V,t}^N, G_{R,t}^N, \sigma \right] - E_0 \left[(xv - p - \kappa) | C, F_V, G_R, \sigma \right] \right\| < 1$$

Therefore, for Δ (and hence η) sufficiently small and $\gamma < \frac{\varepsilon}{3(T+1)}$, we have uniformly over v

$$\left\| \mathcal{V}^N(v, C_t^N, F_{V,t}^N, G_{R,t}^N | \sigma) - \mathcal{V}(v, C, F_V, G_R | \sigma) \right\| < T * \frac{\varepsilon}{3T} + \frac{\varepsilon}{3} + (T+1)\gamma < \varepsilon$$

where the first error term refers to errors that occur in the approximation when equation 65 holds, the second term includes errors accruing in periods after T, and the final term is the expected error from the event when equation 65 fails to hold. □

Now we prove our main approximation result.

Proposition B.3. Consider a SCE (σ, C, F_V, G_R) where $e(C, F_V, G_R) < 1$. For any $\varepsilon > 0$ we can choose $\eta > 0$ and N^* such that for all $N > N^*$ and $\gamma > 0$, σ is an $\varepsilon - BNE$ strategy if the state in the first period of the N-agent game, $\left(C_0^N, F_{V,0}^N, G_{R,0}^N \right)$, satisfies

$$\left\| C_0^N - C \right\| + \left\| F_{V,0}^N - F_V \right\| + \left\| G_{R,0}^N - G_R \right\| < \eta$$

Proof. From the one-step deviation principle, it suffices to consider a deviation by a single agent in a single period. Without loss of generality, let us assume the deviation occurs in period 0 by a bidder that chooses *Enter.* Lemma B.9 implies that for any $\eta, \gamma > 0$ we can choose η sufficiently small that with probability $1 - \gamma$

$$\left\| C_1^N - C \right\| + \left\| F_{V,1}^N - F_V \right\| + \left\| G_{R,1}^N - G_R \right\| < \eta \tag{66}$$

even if a single agent deviates from the SCE strategy in period 1. Lemma B.12 implies that for any $\epsilon > 0$ we can choose $\eta, \gamma > 0$ sufficiently small and N sufficiently large that with probability $1-\gamma$

$$\left\| \mathcal{V}^N\left(v, C_1^N, F_{V,1}^N, G_{R,1}^N \middle| \sigma\right) - \mathcal{V}(v, C, F_V, G_R|\sigma) \right\| < \frac{\epsilon}{4} \tag{67}$$

for $\left(C_1^N, F_{V,1}^N, G_{R,1}^N\right)$ that satisfy equation 66. Equation 67 implies that the effect of the current deviation on future periods is small.

Let $\beta_{dev}(v)$ denote the optimal deviation for type v at $\left(C_0^N, F_{V,0}^N, G_{R,0}^N\right)$. From Lemma A.2 and Assumption 2.5 we have that for any $\delta \geq 0$ and all v that we can choose η sufficiently small that

$$\|\beta_{dev}(v) - \beta(v)\| < \delta v \tag{68}$$

From Lemma B.7 we have for N sufficiently large:

$$\sup_{b,v\in[0,1]} \left\| E_0^N\left[x(b)v - p(b)\left| C_0^N, F_{V,0}^N, G_{R,0}^N\right.\right] - E_0\left[x(b)v - p(b)\left| C_0^N, F_{V,0}^N, G_{R,0}^N\right.\right] \right\| < \frac{\varepsilon}{6}$$

Combining this with Lemma B.11 yields for η sufficiently small and N sufficiently large

$$\left\| E_0^N\left[x(\beta_{dev}(v))v - p(\beta_{dev}(v))|C_1, F_{V,0}^N, G_{R,0}^N\right] - E_0\left[x(b)v - p(b)\,|C, F_V, G_R\right] \right\| < \frac{\varepsilon}{4} \tag{69}$$

Equations 67 and 69 and the fact that the SCE strategy is weakly optimal in the limit game given the SCE aggregate values of C, F_V, and G_R we get

$$\mathcal{V}^N\left(v,\ C_0^N, F_{V,0}^N, G_{R,0}^N \middle| \sigma\right) + \varepsilon \geq E_0^N\left[(x(\beta_{dev}(v))v_i - p(\beta_{dev}(v)) - \kappa) + \right.$$
$$\left. (1 - x(\beta_{dev}(v)))\delta\mathcal{V}^N\left(v_i, C_1^N, F_{V,1}^N, G_{R,1}^N|\sigma\right) \middle| \sigma, C_0^N, F_{V,0}^N, G_{R,0}^N\right]$$

which implies that following the SCE strategy is an ϵ-BNE for $\eta > 0$ sufficiently small. □

Appendix C. ONLINE SUPPLEMENTAL MATERIALS FOR *How Efficient are Decentralized Auction Platforms?* by Aaron L. Bodoh-Creed, Jörn Boehnke, and Brent R. Hickman EMPIRICS: ADDITIONAL DETAILS AND FIGURES

C.1. **Model Tuning Parameters and Additional Figures.** Our estimator from Section 3.2 relied heavily on B-splines to achieve a high degree of flexibility in our functional forms, while maintaining good numerical behavior. One of the benefits of B-splines is their ease of incorporating shape restrictions, many of which can be imposed as simple linear constraints on the parameter values themselves. For example, under the Cox-de Boor recursion formula (with concurrent boundary knots), the only basis functions to attain a non-zero value at the boundaries are $\mathcal{F}_{b1}(\cdot)$ and $\mathcal{F}_{b,P_b+3}(\cdot)$, which both equal one at the upper and lower endpoints, respectively. Therefore, enforcing boundary conditions is equivalent to setting the first and/or last parameter value equal to the known boundary value(s) of the B-spline function, which also cuts down on computational cost by reducing the number of free parameters. Monotonicity is also quite simple: [de Boor, 2001, p.115] showed that a B-spline function $\hat{G}_B(b;\boldsymbol{\alpha}_b)$ 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.

However, before implementing the estimator there remain several free parameters to pin down, the most important of which are the knot vectors $\mathbf{n}_b$, $\mathbf{n}_r$, and $\mathbf{n}_v$ which in turn define the B-spline basis functions. 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 I 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 $I_r = 5$.[42] This gives us a total of 13 parameters for $\hat{G}_B$ and 8 for $\hat{G}_R$. We chose $I_v = 15$ knots at the quantiles of the distribution $\hat{G}_B \circ \hat{\beta}$, which is known from Stage I. We chose $I_v > I_b$ because $\hat{F}_V$ must conform to the nuances induced by all first-stage parameters in order to accurately represent the implied private value distribution. We find that these choices provide a good fit to the data and that adding more parameters renders little benefit.[43]

Figure 9 displays plots of 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 9 also provides intuition about the nature of the B-spline functions we use to fit the variables of our model.

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

[43] 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.

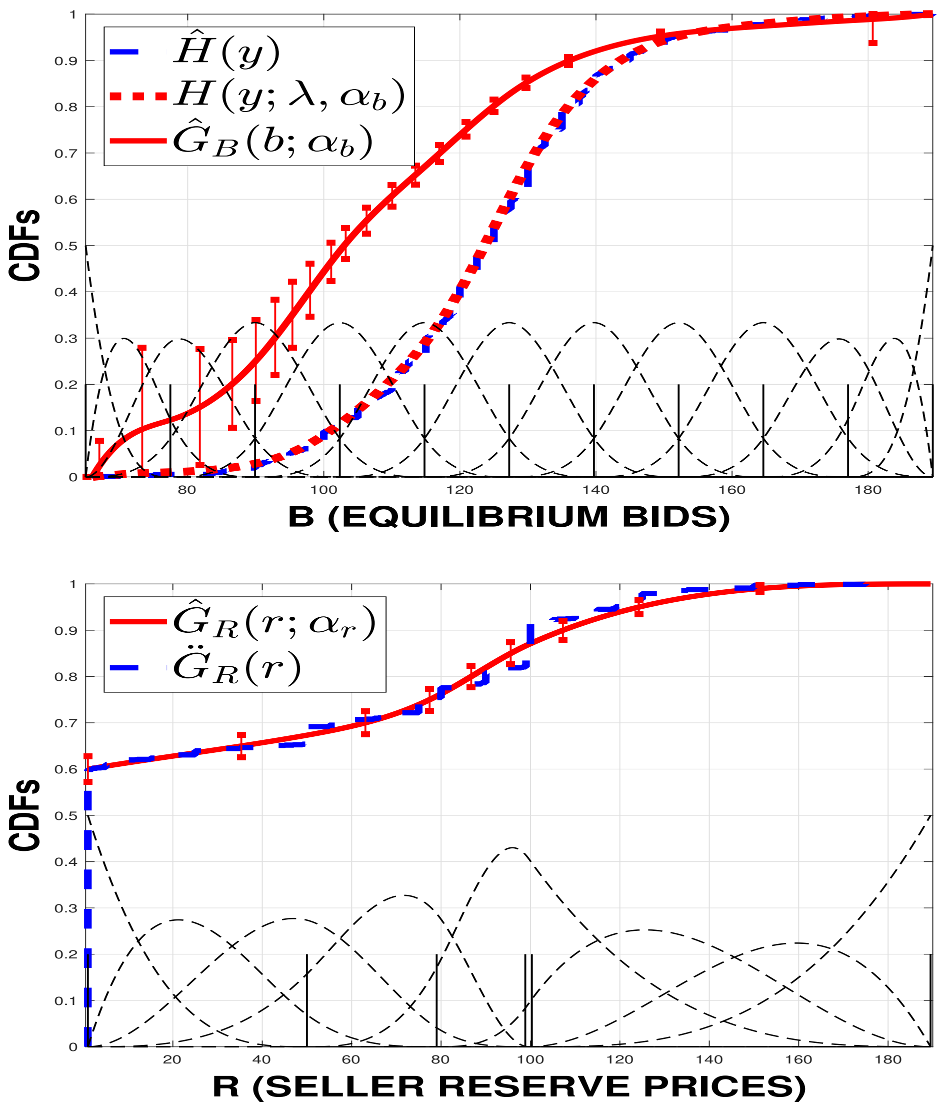

FIGURE 9. Stage I Estimates

Since both panels describe observable variables, we have included the empirical CDFs as well. A comparison between the empirical CDFs and the B-spline fit shows a very close correspondence. To illustrate the underlying components of our B-spline functions, we have included the locations of the knots and the basis functions in the plot as well. The knot locations are described by the thin, vertical, solid lines extending at the bottom of each panel. The families of basis functions are drawn at the bottom of each panel in thin, dashed lines.

APPENDIX D. ONLINE SUPPLEMENTAL MATERIALS FOR
How Efficient are Decentralized Auction Platforms?
BY AARON L. BODOH-CREED, JÖRN BOEHNKE, AND BRENT R. HICKMAN

ADDITIONAL COUNTERFACTUAL RESULTS

D.1. **Revenue Counterfactuals.** One of the main findings of Section 5.1.3 is that the revenue of the sellers decreases as the efficiency of the within-period allocation increases. From the Revenue Equivalence Theorem, we know that the revenue is determined by (1) the allocation rule (i.e., $\chi_u(b)$) and (2) the steady-state distribution of types. In order to tease apart these effects, we recomputed our revenue and welfare statistics holding the measure and type distribution of the buyers fixed at the status-quo values, which in turn means that all of the efficiency and revenue changes are driven by the allocation rule alone.

Unlike the analysis of Section 5.1.3, all of the analysis in this section is out-of-equilibrium since we are not accounting for how altering the allocation function changes the steady-state measure and type distribution of the buyers. It is somewhat misleading to claim that we are isolating the effect of the allocation mechanism since the allocation mechanism is a determinant of the steady-state distribution of types. The most natural interpretation of our exercise is that we are studying how our predictions would differ had we used a mis-specified model that assumed the measure and type-distribution of the agents are exogenous objects. As we will see, assuming the population of buyers is exogenous will lead us to over-estimate the efficiency gains from centralization. More surprisingly, we will find the mis-specified model would predict a sharp *increase* in auction revenues, which is the opposite of what the true model predicts. The primary take-away from our exercise is that it is absolutely crucial to account for the endogenous changes to the population of buyers when producing counterfactuals.

We now step through the changes to the endogenous objects (i.e., $\chi_u(b)$, $\mathcal{V}(v)$, and $\beta(v)$) when evaluating the centralization counterfactual using our mis-specified model. The left pane of Figure 10 plots $\chi_1(b)$ and $\chi_8(b)$ for $\delta = 0.88$ under the status-quo steady-state type distribution.[44] Recall that the value function can be described as a function of $\chi_u(b)$ using the Envelope Theorem as per equation 37, reproduced here for clarity:

$$\frac{\partial \mathcal{V}(v)}{\partial v} = \sum_{\tau=t}^{\infty} \delta^{\tau-t}(1-\chi_u(\beta(v)))^{\tau-t}\chi_u(\beta(v))$$

The derivative of the value function is shown in the right pane of Figure 10. The changes to the allocation function make it less likely that low-value buyers win and has the opposite effect for high-value buyers. This causes the derivative of the value function to drop precipitously for low value buyers and rise for high-value buyers.

The net effect on the value-function is displayed in the left pane of Figure 11. As one can see, the value function is uniformly *lower* when $u = 8$ than when $u = 1$, although the difference is small for the highest value buyers. This contrasts with the results of Section 5.1.3 wherein the value function for the agents that remain in the market increases as u

[44]The results are qualitatively identical for other choices of δ.

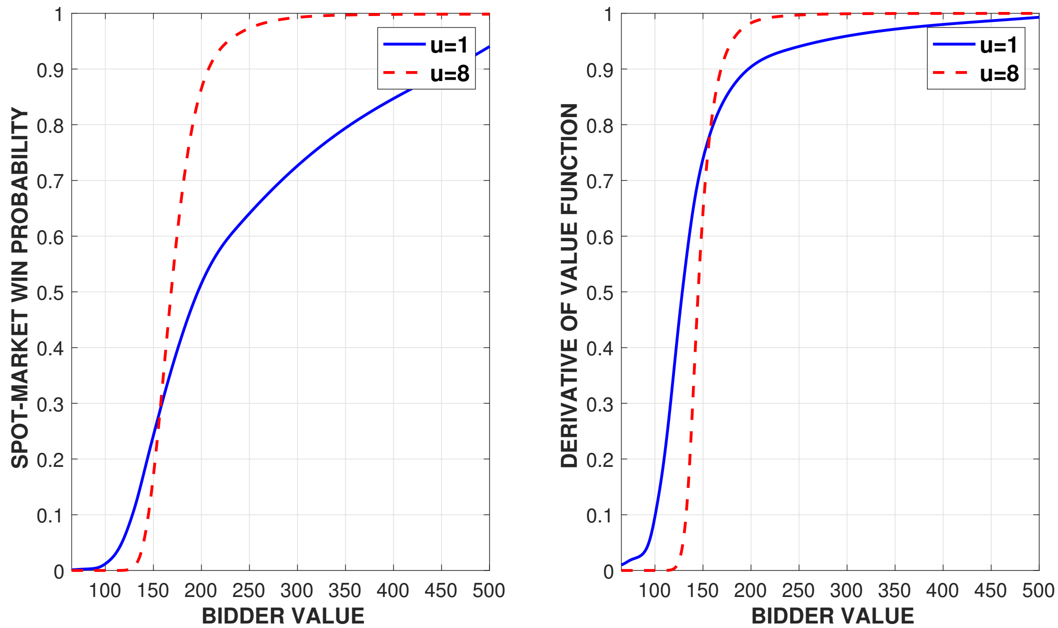

Figure 10. Out-of-Equilibrium Probability of Winning

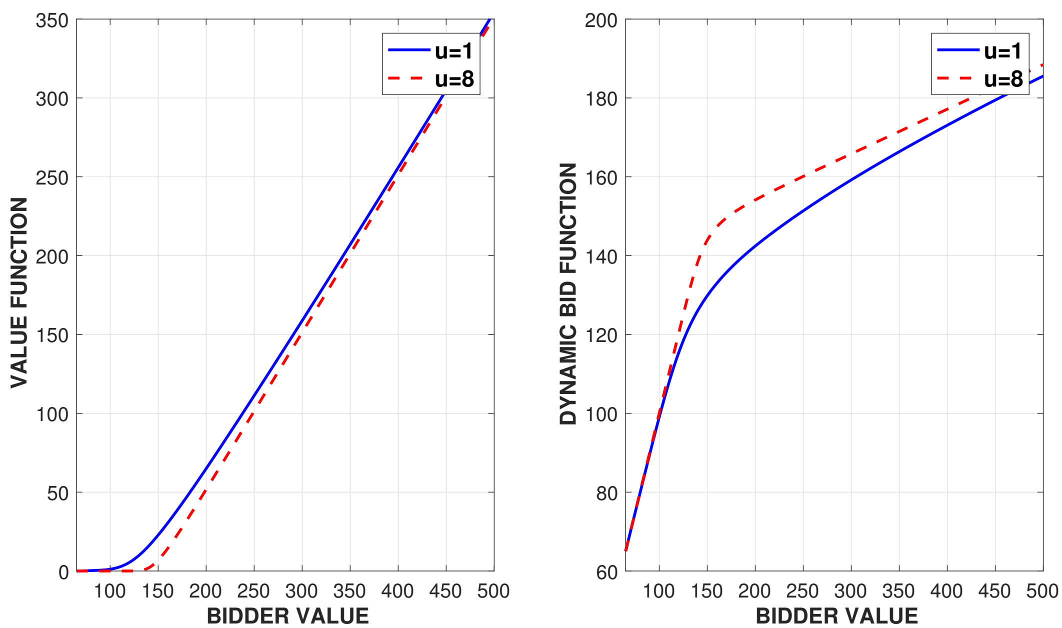

Figure 11. Out-of-Equilibrium Bids

rises. The decrease in the value function has the corresponding effect of increasing the agent's bids, which are displayed in the right pane of Figure 11.

In an SCE with $u = 8$, high value buyers exit the platform quickly, whereas there is a large population of high value buyers in the status quo F_V. If we treat the status quo

TABLE 6. Mis-specified Counterfactual Efficiency Ratios $\mathcal{E}_{u,\delta}$

# Units Per Listing	**Discount Factor $\delta =$** 0.75	0.80	0.86	0.88	0.92	0.95	0.98
1	0.89	0.88	0.87	0.87	0.85	0.84	0.82
2	0.94	0.94	0.93	0.93	0.92	0.92	0.91
4	0.97	0.97	0.96	0.96	0.96	0.96	0.95
8	0.98	0.98	0.98	0.98	0.98	0.98	0.98
Lottery	0.58	0.55	0.50	0.47	0.41	0.35	0.25

TABLE 7. Mis-specified 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	\$120.83	\$120.77	\$120.68	\$120.63	\$120.52	\$120.37	\$120.06
2	\$129.33	\$130.00	\$131.17	\$131.91	\$133.80	\$136.90	\$145.95
4	\$134.11	\$135.52	\$138.03	\$139.59	\$143.67	\$150.49	\$171.54
8	\$137.11	\$139.23	\$143.00	\$145.36	\$151.60	\$162.29	\$197.46

measure F_V as fixed and exogenous, then the buyers face increasingly intense competition for the good as u increases since it becomes increasingly unlikely that any bidder with a value outside the top quintile of the value distribution will win the item. Even for buyers with those high values, the expected price increases because of the intense competition from other high value buyers. These competition effects depress the value function for all types of buyers.

The efficiency ratios and revenues of our out-of-equilibrium counterfactual are higher due to two effects. First, the steady-state distribution has a large population of high valuation buyers relative to the equilibrium type distribution for $u = 8$. In equilibrium, these high-value agents are quickly allocated a good and leave the economy, and so the population of these agents is relatively small in equilibrium. Second, since the bids are higher, they are more likely to surpass the starting price, although this effect is far less important than the first effect. The higher bids also has the effect of increasing the auction revenue, whereas the lower bids in the $u = 8$ SCE depress the revenues. The out-of-equilibrium efficiency ratios are displayed in Table 6, while the revenues are displayed in Table 7.

D.2. **Participation Cost Counterfactuals.** We would like to highlight the effect of centralization on the participation costs of the buyers through two different metrics. The first metric is the flow of average lifetime participation costs (FALPC) paid by the agents. At

Table 8. Effects of Centralization on FALPC

# Units	Discount Factor $\delta =$						
Per Listing	**0.75**	**0.80**	**0.86**	**0.88**	**0.92**	**0.95**	**0.98**
1	\$0.52	\$0.52	\$0.53	\$0.53	\$0.53	\$0.53	\$0.54
2	\$0.37	\$0.37	\$0.37	\$0.38	\$0.38	\$0.38	\$0.38
4	\$0.27	\$0.27	\$0.27	\$0.27	\$0.27	\$0.28	\$0.29
8	\$0.20	\$0.20	\$0.20	\$0.20	\$0.21	\$0.21	\$0.23

the point an agent of type v bids for the first time, the expected lifetime participation costs paid by the agent is equal to $\kappa/\chi_u(\beta(v))$. The FALPC measures the total participation costs paid by the set of buyers entering the game in the current period. A measure μ of such buyers with types distributed as per $t_V(v)$ enter each period, meaning that the FALPC can be written as:

$$\mathcal{FALPC} = \mu \int_0^\infty \frac{\kappa}{\chi_u(\beta(v))} t_{V,\delta}(v) dv$$

By focusing on the newly entering bidders, we are measuring the increase in the total expected participation costs in the current period, hence our choice to refer to this statistic as a "flow." Since the distribution of entering bidder types is a structural primitive, the only effect of changing the spot-market mechanism is to change $\chi_u(b)$. The FALPC values are summarized in Table 8. Because the FALPC is not based on agent value, δ has very little effect on the realized values. On the other hand, u has a strong influence through its effect in $\chi_u(b)$.

Across all of the discount factors, centralizing from $u = 1$ to $u = 8$ results in a roughly 60% decline in FALPC. There are two channels for the reduction in participation costs. First, because the good is allocated more efficiently when $u = 8$, high-valuation buyers are matched to the good more quickly. Second, and more importantly, since low-valuation buyers are less likely to win the good each period, these agents leave the game and do not incur participation costs in the first place. Inducing these buyer-types to exit has a particularly strong effect on participation costs since these low-valuation buyers stay in the market for many periods before winning when $u = 1$.

Our second metric, the stock of average lifetime participation costs (SALPC), assesses the expected lifetime participation costs of he agents that are present in the "stock" of agents characterized by the steady-state type distribution, $f_V(v)$. The equation for the SALPC is identical to equation 70 once we replace the measure and type distribution of the "flow" of entering buyers with the steady-state measure and distribution of agent types

TABLE 9. Effects of Centralization on SALPC

# 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

that characterize the "stock" of agents present in each period:

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

The SALPC is described in Table 9.

There are two effects at work. First, when markets centralize, the total participation costs paid by a participant before winning an item goes slightly up on average. For example, when $\delta = 0.88$ and $u = 1$,[45] the participants paid on average $8.98 each over their lifetimes in the market. When $\delta = 0.88$ and $u = 8$, the participants paid on average $9.50 over the course of their participation in the market. The average participation cost per bidder is pushed up by the fact that lower value agents must wait even longer (on average) before winning an item and exiting the market.

The larger effect is that fewer buyers participate in the market when u increases. The buyer to seller ratio is 7.75 when $\delta = 0.88$ and $u = 1$, while the ratio is only 2.92 when $\delta = 0.88$ and $u = 8$. The SALPC is the product of the average per-bidder cost and the ratio of buyers to sellers. As seen in Table 9, the SALPC drops by roughly 60% as u moves from 1 to 8 for the $\delta = 0.88$ case.

Obviously the FALPC values are much lower than the corresponding SALPC values. The gap reflects the differences between the structural primitives describing the agents that enter each period and the steady-state distribution. In short, because the steady-state includes a larger population of low value buyers that must wait many periods to win, the SALPC will disproportionately reflect the high lifetime participation costs incurred by these low value buyers.

We view the FALPC and SALPC as complementary metrics. The SALPC metric better reflects the population of buyers present at a given time, which makes SALPC the better metric of the participation costs as perceived by the population of agents typically present on the platform. If one were interested in the net welfare effect of participation costs, one

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

could use FALPC as a metric of per-period welfare loss and compute the net present value of future welfare losses. Given the daily rate at which FALPC is measured, the appropriate discount factor ought to be very close to 1, which would result in net welfare losses of the same order of magnitude as the SALPC calculations presented in Table 9.

To place these results in the context of our finite model, one needs to recall that we normalize the measure of sellers to 1, which means one can interpret the costs above in terms of the cost per auction. For example, centralizing from $u = 1$ to $u = 8$ causes the FALPC to drop from \$0.53 to \$0.21 per auction when $\delta = 0.95$. Computing the net present value using this same discount factor, we find that centralization reduces the net present value of the participation costs by \$6.40 per auction, per day. Again, since this is a daily discount factor, we view $\delta = 0.95$ as an extremely low estimate of the appropriate time discount factor for computing the net present value. To find the total cost incurred in the finite setting each day, we need to "de-normalize" the measure of sellers by multiplying the per-auction statistics by the 11.25 auctions per day that we observe in our data. Given this average number of auctions per day, the total welfare gain from the reduction of the participation costs per day is \$72 per day.[46]

D.3. **Seller Incentives.** As has been regularly noted about the eBay marketplace, sellers tend to choose very low starting prices. In our data, almost 60% of the starting prices are set at the lowest possible value of \$0.99. It is easy to see that such a price is not optimal - a single seller could improve his profits if he set a starting price equal to $\underline{v}$, the lowest possible bid in the auction.[47] From the perspective of a single seller, choosing the optimal starting price involves the same reasoning as in the classic optimal auctions literature: a high starting price can increase the final price paid by the winner, but it also risks that the good will go unsold.

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

$$max_{r \geq 0}\ Pr\left[B^{(1)} \geq r\right] E\left[max\{r, B_M\} | B^{(1)} \geq r\right] \tag{70}$$

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

[46] One might have compared the welfare losses from the participation costs with the welfare gains created by the market. Unfortunately, we do not believe such an exercise is credible since the consumer surplus is highly dependent on the choice of δ whereas the participation costs are insensitive to δ.

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

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

The benefits from optimally choosing the starting price are small because each seller is matched with 7.96 bidders in expectation, which means that the competition between bidders is intense. Bulow and Klemperer [1996] show that in a static auction setting choosing the starting price optimally pales in comparison to adding a single extra bidder to the market.[48] With almost 8 bidders on average already participating, it should not be surprising that there is little room left for optimizing the starting price to have a significant effect on auction revenues.

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