Selling Mechanisms for Perishable Goods:

An Empirical Analysis of an Online Resale Market for Event Tickets

Caio Waisman*

February 25, 2021

Abstract

This paper assesses the value of the availability of menus of different selling mechanisms to agents in an online platform in the context of scarce perishable goods. By analyzing the choice between auctions and posted prices in the context of National Football League tickets offered on eBay, it estimates a structural model in which heterogeneous, forward-looking sellers optimally choose which selling mechanism to use and its features. Counterfactual results suggest that sellers would experience an average 87.37% decrease in expected revenues if auctions were removed and just a 4.34% decrease if posted prices were. In turn, buyers would benefit from an auction-only platform since the expected number of transactions would increase and expected transaction prices would decrease. These results suggest that while sellers benefit from menus of different selling mechanisms, the same does not hold for buyers. Thus, the implications for a platform, which should take into account both sides of the market, are ambiguous.

JEL: C57, L11, L81, L83, M31

Keywords: mechanism choice, auctions, posted prices, perishable goods, platform design

*Marketing Department, Kellogg School of Management, Northwestern University. E-mail: caio.waisman@kellogg.northwestern.edu.

1 Introduction

The choice of which selling mechanism to use is one of the most fundamental decisions a seller can make. Theoretical work on this topic has been extensively conducted in the fields of economics, marketing, finance, operations research, and computer science. More recently, the advent of technology-based marketplaces that bring buyers and sellers together to facilitate trade has made it easier for different selling mechanisms to be, often concurrently, available and employed in practice. These platforms have made use of a wide array of different mechanisms over time, most notably auctions and posted prices, which illustrates the relevance of mechanism choice in practice: TaskRabbit and Prosper.com began their operations as auction-based platforms, but abandoned auctions and focused on improving matching procedures; eBay began as an auction-only market but now hosts several different selling mechanisms; and Upwork and Freelancer still rely heavily on procurement-like mechanisms to match workers to employers for specific tasks. This raises a question: what is the value, if any, of the availability of menus of different selling mechanisms to agents in these markets? The main goal of this paper is to assess what this value is for buyers and especially sellers on an online platform.

To make this assessment, it is necessary to compare market outcomes from an environment where agents have access to a menu of different selling mechanisms with outcomes from alternative environments where agents have access to only one such mechanism. I make this comparison using data on National Football League (NFL) tickets listed on eBay. This is a favorable setting to perform this comparison because eBay provides sellers with a menu of different selling mechanisms from which they can choose how to list their goods. I particularly focus on the tradeoff between auctions and posted prices because these are the main mechanisms available at eBay.

Since I do not observe an environment where eBay offers only one selling mechanism, I have to estimate a structural model so I can perform the aforementioned comparison. One of the benefits of focusing on event tickets is that they, like airline seats, hotel rooms, and online advertising spots, are perishable goods: they need to be consumed before or at a deadline. The existence of a deadline shifts the incentives of agents to choose specific selling mechanisms, which naturally yields a framework upon which a parsimonious model can be built. In addition, data from this setting display features that are useful for the purposes of estimation such as variation in mechanism choice within the same seller-tickets pair over time, which is helpful to separate the drivers of mechanism choice from seller and product heterogeneity. However, it is also important to recognize that deadlines introduce additional complexity to the problem of how to sell goods because the environment in which they are traded is inherently nonstationary. Nevertheless, this same complexity makes the improvement of markets for perishable goods a continuing effort, with the use of different selling mechanisms being frequently suggested as a way to achieve progress. Thus, this paper also takes a step in this direction, even if small.

The proposed structural model is specified to capture the empirical patterns observed in the data and institutional details of eBay while maintaining a set of simplifying assumptions for the sake of tractability. Its main feature is that sellers are forward-looking and heterogeneous, where heterogeneity is with respect to their outside options and listing cost parameters, which can also be interpreted as an inconvenience or monitoring disutility. This heterogeneity is required to fit the data and is also directly connected to observed differences in mechanism choice across sellers. All else constants, sellers with higher listing costs are more likely to list their tickets with posted prices because, unlike auctions, they do not have an end date. Hence, sellers could use posted prices and just leave the tickets available on the platform without the need to closely monitor and relist them. Outside options, in turn, capture the patterns for seller exit. Finally, the demand side is more simplified and, following a relevant subset of the revenue management literature, taken as exogenous.

Based on the estimates of this model, I then conduct counterfactual exercises in which the menu of mechanisms is altered. Motivated by recent trends and events in online platforms, I first investigate what would happen if auctions were eliminated from the platform. The effects are highly adverse to sellers, leading to decreases in both the probability of sales and expected transaction prices and an ultimate decrease of more than 87% in overall expected revenues. On the other hand, creating an auction-only platform would have milder negative effects on sellers as an increase in the probability of sales partially offsets the decrease in expected transaction prices, implying an overall decrease in expected revenues of less than 5%. Importantly, buyers would benefit from such a platform as not only it makes purchases more likely but also at lower prices. This creates a potential tension for an online platform that has to cater to both sides of the market in order to sustain market operation.

The finding that auctions are preferred by both sides of the market is surprising given the aforementioned recent transition of online platforms away from this selling mechanism. There are at least two reasons for why this could be the case. First, this study is confined to a market for perishable goods, while platforms that abandoned auctions and studies that analyzed eBay's transition to posted prices focused on goods and services of a different nature. Second, this result was obtained under several assumptions that were kept to maintain tractability of the model. Returning to this question with richer data and a more intricate model is, therefore, a valuable exercise for future research.

Overall, the main contributions of this paper are twofold. The first contribution is a detailed description of a market for perishable goods with mechanism choice, which, in turn, displays two novelties compared to other studies. First, the ability to track the same seller-tickets over times allows me to show how, when and whether sellers choose and switch between different mechanisms. Second, I have access to clickstream data that enable me to observe all the buyers who click on ticket listings on a given day even if they do not make purchases or submit bids, which I can then employ as a measure of potential demand. This information on potential demand is, to my knowledge, novel, and therefore describing it can be of independent value for researchers and practitioners specifically interested in markets for perishable goods.

The second main contribution is assessing the value of providing a menu with different selling mechanisms to agents in a market for perishable goods while accounting for dynamics. As I detail below, most existing papers compare market outcomes when either auctions or posted prices are available, but not offered concurrently. Studies that addressed dynamic features of market for perishable goods, in turn, did so only in environments where posted prices were the sole mechanism in use. This paper innovates by analyzing the concurrent use of selling mechanisms in this specific type of market, whose design, as aforementioned, is a continuing effort.

The remainder of the paper proceeds as follows. First, I review the existing literature to which this paper is related. Next I describe the date used in this study and document the main empirical patterns they display. I then propose an empirical model that aims to rationalize sellers' choices and discuss how this model's primitives can be estimated, which is followed by the presentation of this model's estimates and counterfactual results. Finally, I conclude the paper by summarizing the key findings and outlining directions for future research.

1.1 Related literature

The theory of mechanism choice and, more broadly, mechanism design, has been extensively studied. The specific tradeoff between and simultaneous use of auctions and posted prices have also received considerable attention due to the popularity of these mechanisms, with notable contributions by Wang (1993), Kultti (1999), Julien et al. (2001), Ziegler and Lazear (2003), Etzion et al. (2006), Sun (2008), Eeckhout and Kircher (2010), Etzion and Moore (2013), Hummel (2015), Selcuk (2017), Maslov (2020a), Maslov (2020b), and Zhang (2020). The problem of how to sell a good before a deadline has been widely addressed in the revenue management literature as the textbook by Talluri and van Ryzin (2004) attests. Even though this literature traditionally studies the optimal policy given a mechanism, the simultaneous use of auctions and posted prices has recently been addressed, for example, by Caldentey and Vulcano (2007), while a recent study by Board and Skrzypacz (2016) combines mechanism design with revenue management.

Empirical studies addressing the choice between auctions and posted prices were facilitated by the availability of data from online platforms. A stream of the literature focused on the platform's choice of which mechanism to offer to its users. For instance, both Wei and Lin (2017) and Huang (2020) studied Prosper.com's decision to switch from auctions to posted prices. Since in this context only one mechanism existed at any given point in time, these studies did not have to consider the problem of how buyers and sellers choose between different mechanisms. However, they also did not compare outcomes from environments where a single mechanism could be employed with those from an environment in which both mechanisms coexisted.

This paper relates more closely to studies that focus on markets where both mechanisms can be concurrently employed. Early contributions were made by Zeithammer and Liu (2006), who found that observed and unobserved seller heterogeneity are the main drivers of mechanism choice, and Hammond (2010), who documented that auctions are more likely to convert but posted prices yield higher transaction prices, a dichotomy also found in this and other papers. Seller behavior and the recent preference towards posted prices were further addressed by Einav et al. (2015) and Einav et al. (2018). Einav et al. (2015) made use of a large data set to study seller strategies on eBay, in particular episodes in which sellers offer similar products with different mechanisms or prices. While this strategy could be valuable to sellers as a tool to learn about demand, the extent to which it can be used to sell a perishable good is limited by the deadline. Furthermore, few sellers in my data offer several ticket bundles for the same game, and the practice of concurrently offering the same set of tickets with different mechanisms is disincentivized by eBay. In turn, Einav et al. (2018) focused on the secular trend of decrease in use of auctions, indicating that this is as a result of posted prices being more convenient despite their more limited potential for price discovery. This paper's finding that posted prices become relatively more popular as deadlines approach is somewhat reminiscent of their convenience incentive. Nevertheless, this paper focuses on the strategic choice between mechanisms to sell the exact same good, while Einav et al. (2018) focus more on the functioning of the market as a whole. More recently, Coey et al. (2020) proposed an equilibrium search model in which buyers have private deadlines that lead auctions to become discount mechanisms. This paper's environment is very different from theirs, however, because deadlines not only are public but also shared by buyers and sellers.

Other papers introduced structural models to explain mechanism choice, with an initial contribution by Hammond (2013). The studies to which this paper is most closely related are Sweeting (2013) and Bauner (2015), since both introduced structural models of mechanism choice in the context of perishable goods, namely Major League Baseball (MLB) tickets offered on eBay. Both these studies also accounted for hybrid, buy-it now auctions, which I ignore, but they treated the sellers' problem as a static one. Furthermore, these three studies relied on seller-specific preferences towards specific mechanisms to explain mechanism choice. In turn, I leverage data on repeated choices at the seller-tickets level to model the sellers' problem as a dynamic mechanism choice one, using their payoff maximizing decisions to understand mechanism choice without the need for seller-specific inherent preferences towards particular mechanisms. While other studies also modeled sellers' dynamic behavior when offering event tickets, such as Sweeting (2012), Lee et al. (2012), Sweeting and Sweeney (2015) and Sweeting (2015), they focused solely on pricing. Hence, this paper makes a contribution by empirically addressing dynamic mechanism choice in the context of a perishable good.

Finally, the topic of event ticket resale has received considerable attention as the survey by Courty (2003) demonstrates. The usual focus of this literature is on arbitrage, with a recent empirical contribution by Leslie and Sorensen (2014). More related to this paper is the contribution by Bhave and Budish (2018), which leveraged an experiment to show how the use of auctions in the primary market for tickets can be an effective tool to mitigate arbitrage opportunities. While this paper's setting is a secondary market for event tickets, it does not explicitly address issues related to arbitrage or resale.

2 Data and descriptive analysis

The goal of this paper is to assess the value of menus of different selling mechanisms to agents on an online platform. I focus more specifically on sellers, who are the agents that actively choose the selling mechanism used to list their goods, using data from eBay on NFL tickets. I now describe these data.

My goals in this description are threefold. First, since this setting is a market for perishable goods, I document how the market evolves dynamically, especially as a function of proximity to the deadline, which is game day. This is important because the empirical model is constructed around this particular dynamic structure, and therefore demonstrating that agents behave in a way consistent with these dynamic considerations is crucial. Second, for the sake of tractability some of the modeling decisions I make are not supported by the data. It is my intention to be transparent about such cases and to justify why they are present. Finally, documenting how a market for perishable goods where different selling mechanisms are concurrently available evolves over time can be of independent interest. In particular, the clickstream data on the buyer side are, to my knowledge, novel. Therefore, describing when and how buyers click on existing listings for tickets over time is an additional contribution of this paper.

I begin by listing which variables the data contain, followed by a brief summary of the final sample I use to estimate the structural model. I then look at the supply and demand sides of this market separately, focusing on how each side's decisions change as a function of distance to the deadline. Finally, I summarize how the market as a whole evolves over time.

2.1 Observed variables

The bulk of the data used in this study comes from eBay. For all listings of NFL tickets created on eBay between January of 2013 and February of 2014, I observe: when the listing was created and when it ended; its format (auction, hybrid auction, or posted price with or without a bargaining option); for auctions, the start, reserve and buy-it-now prices;[1] posted prices for fixed-price listings and all their changes; the duration of auctions; the

[1] On eBay, start prices are the price at which bidding starts, acting effectively as a public reserve price, while reserve prices are private and only disclosed if bids are submitted. Since reserve prices are rarely binding in my data, I will ignore them and use only start prices, referring to them as reserve prices instead.

title and subtitle chosen by sellers; the product category; whether the listing was sold; the number of bundles and tickets per bundle; the location of the seats at the section and row level; the game to which the tickets corresponded; and the identity of sellers. Since the vast majority of tickets in the data are electronic, I abstract from shipping and delivery considerations. Regarding sellers, I further observe their user name and their feedback score and percentage rating over time. The score is computed in the following way: a seller receives one point for each positive rating, no points for each neutral rating, and loses one point for each negative rating. The percentage rating is simply the fraction of all ratings received by a seller that were positive.

In addition to these variables, the data also contain transaction prices, bargaining offers, and bids. Furthermore, I also observe a measure of potential demand. This measure is different from what is commonly observed in marketing and industrial organization studies. For example, scanner data usually only contain information regarding transactions; these data do not record cases of consumers who entered a store but did not purchase anything. Information about shopping visits with no purchases is directly analogous to the measure of potential demand I employ. More specifically, for each listing-day pair, I observe all users who clicked to view this listing's detail page and how many times each user clicked. Users are classified based on their IP number. Since it is not possible to know whether the same individual clicked on a given listing with different IP numbers unless this user was logged in, each number is treated as a different potential buyer. Based on this variable, I am able to compute a measure of relative scarcity I refer to as market tightness: for each game-day pair, it consists of the number of different users who were observed clicking on at least one listing for tickets divided by the number of available listings for tickets. Despite being noisy, I use this measure to illustrate how the market evolves over time, particularly as the game approaches. It also plays a key role in the structural model because it directly influences the number of buyers who arrive to listings. However, it is important to note that I do not observe potential buyers who query tickets for a given game and, upon seeing the results from this query, choose not to click on any of the available listings.

Finally, I use the face value of the tickets, that is, their original prices in the primary market, to compare monetary amounts in the descriptive analysis and parametrize sellers' outside options in the model. To recover ticket prices I first made use of the Wayback Machine through the Internet Archive website to access each team's web page in mid-2013. When pricing schedules were not available through this resource I contacted each team separately to try to obtain past prices directly from them. At the end of this process

I was able to recover prices for all teams except the then named Washington Redskins, which then led me to remove all ticket listings associated with their home games from the data. These prices referred to the individual ticket price charged in season passes, and were matched to each listing based on the location of the tickets.

2.2 Overview of the final sample

The structural model focuses on the sellers' dynamic choice of which selling mechanism to employ when they list their tickets. Hence, the relevant unit of observation is a seller-tickets pair, and estimation is performed by matching the predicted sequence of listings the seller chooses for a given set of tickets implied by the model to what I actually observe in the data. I refer to this sequence of listings by the same seller for the same set of tickets as a chain. An example of a chain observed in the data is given in Table 1. In this example, the seller first listed the tickets ten days before the game as a posted price, charging 125 dollars for them. One day later, the seller relisted the tickets as a three-day auction with a reserve price of 99 dollars. After this auction was unsuccessful, the seller relisted the tickets once again as a posted price but only charging 89 dollars for them. A sale was made two days before the game.

Table 1: Example of a chain

Chain	Seller	Game day	Created	Ended	Sold	Format	Buy/start price
4	32	12/29	12/19	12/20	No	Posted price	125
4	32	12/29	12/20	12/23	No	Auction	99
4	32	12/29	12/23	12/27	Yes	Posted price	89

It is important to note that the data originally were not structured to keep track of the same item across time, only of different offerings. In other words, recovering each row of Table 1 was straightforward, but determining that the different rows corresponded to the same set of tickets was not. Hence, the chains were not directly available; I created them manually through a process that is described in detail in Appendix A.

Key quantities from the final sample are displayed in Table 2. The final sample consists of 19,174 pairs of tickets that were never offered in bundles that included other tickets and that were never offered in more than one listing at the same time. These 19,174 pairs were offered by 8,081 sellers, encompassed 28,257 different listings, and corresponded to 245 regular season games.[2] Thus, each set of tickets was listed on average 1.5 times, and each seller offered on average 2.4 pairs of tickets.

Table 2: Snapshot of final sample

Variable	Quantity	% sold
Listings	28,257	37.36
Sets of tickets	19,174	55.06
Sellers	8,081	–
Games	245	–

Table 3 shows how the 28,257 listings from Table 2 are distributed across mechanisms. Since my focus is on auctions and posted prices, hybrid auctions are treated as simple auctions unless the buy price was accepted or the listing was created on game day, in which case the listing was considered a posted price, and bargaining-enabled listings are always treated as regular posted prices. A more detailed analysis incorporating these four possible mechanisms is given in Appendix B, where I argue that this simplification is inconsequential. The majority of listings are offered through auctions, which are more likely to be successful than posted prices. In addition, the systematic use of both mechanisms is a first indication that sellers probably benefit from the availability of menus of different selling mechanisms from which they can choose.

[2]These are all eligible games because there are 256 regular season games in total, but Washington Redskins (8) and International Series (3) games were excluded.

Table 3: Distribution of listings across mechanisms

Type	Quantity	% of total	% sold
Auctions	19,082	67.53	38.93
Posted prices	9,175	32.47	34.09
Total	28,257		37.36

Notes: Hybrid auctions are included within auctions with the exception of those that were sold via the buy-it-now option or created on game day, and bargaining-enabled posted prices are included as usual posted price listings.

Table 4 further characterizes the types of chains in the data by displaying whether and how tickets that went unsold were relisted. A little less than three fourths of all sets of tickets are only made available once, in part due to a higher conversion rate: almost 60% of these tickets are sold, while just a little more than 41% of relisted tickets were ever sold. Among the sets that get relisted, the majority are always made available through the same format, which suggests that sellers possibly have inherent preferences that lead them to consistently favor specific mechanisms, and once more there is some indication that auctions are more likely to be successful by comparing the second and third rows.

Table 4: Types of chains

Type	Quantity	% sold	% of all chains	% of all listings
Single-listing	14,162	59.9	73.86	50.12
Multi-listing, always auctions	2,731	39.62	14.24	27.03
Multi-listing, always posted prices	1,071	36.23	5.59	9.47
Multi-listing, mechanism changes	1,210	49.92	6.31	13.38
Total	19,174	55.06		

Notes: Hybrid auctions are included within auctions with the exception of those that were sold via the buy-it-now option or created on game day, and bargaining-enabled posted prices are included as usual posted price listings. A mechanism change is defined as going from any auction to any posted price or vice-versa.

I now separately describe the patterns of mechanism choice for the two types of chains displayed in Table 4, namely single-listing and multi-listing chains. Table 5 displays which mechanisms were used for the single-listing chains along with their rates of conversion. Almost two thirds of such chains were auctions, and roughly two thirds of these auctions were successful, while the conversion rate of posted prices was a little less than 47%. The sets of tickets that went unsold subsequently exited the market without being relisted.

Table 5: Single-listing chains and mechanisms

Type	Quantity	% of total	% sold
Auctions	9,268	65.44	66.72
Posted prices	4,894	34.56	46.98
Total	14,162		50.12

Notes: Hybrid auctions are included within auctions with the exception of those that were sold via the buy-it-now option or created on game day, and bargaining-enabled posted prices are included as usual posted price listings.

Having described mechanism choice in the context of single-listing chains, I now proceed to illustrate it for multi-listing chains. In particular, I focus on the transition patterns between mechanisms when tickets go unsold and are subsequently relisted, which are given in Table 6. The rows show which format was chosen before and the columns indicate the sellers' new choice conditional on the tickets being offered again. There is considerable persistence in sellers' mechanism choice, which once again suggests that sellers might have preferences that lead toward specific mechanisms. However, it is also interesting to note that posted prices are more likely to be relisted as auctions than the converse: 87.42% of relisted auctions re-entered the market as auctions, but only 70.25% of relisted posted prices re-entered the market as such.

Table 6: Relisting transition probabilities (in %) across mechanisms

	To auctions	To posted prices
From auctions	87.42	12.58
From posted prices	29.75	70.25
Total	70.54	29.46

Notes: Hybrid auctions are included within auctions with the exception of those that were sold via the buy-it-now option or created on game day, and bargaining-enabled posted prices are included as usual posted price listings. Probabilities refer to going from the row mechanism to the column mechanism.

Another relevant dimension of the analysis is the type of sellers. It can be expected that a casual season pass owner who lists only one set of tickets due to an idiosyncratic shock behaves differently than a broker who offers several sets of tickets. To verify whether this is the case, I classify sellers into three types: “small" (lists at most two sets of tickets for only one home team), “medium" (lists more than two sets of tickets but for only one home team), and “large" (list sets of tickets for more than one home team). Table 7 breaks down the overall sample characteristics by these types. The majority of sellers in the sample are small, followed by medium and then large. Finally, it is interesting to note that the fraction of sellers that make use of both auctions and posted prices is very similar across seller types, which suggests an unobserved seller-specific driver for mechanism choice. When estimating the model, my goal will be to recover the distribution from which sellers draw their parameters, and I allow the parameters of this distribution to vary across the three different seller types defined here.

Table 7: Types of sellers

Type	Quantity	% of sellers	% of listings	% of chains	% that uses both mechanisms
Small	5,214	64.52	29.36	31.46	6.79
Medium	2,550	31.56	56.57	54.63	6.43
Large	317	3.92	14.07	13.91	6.31
Total	8,081				6.66

Notes: Seller types follow the definitions given in Section 2.2.

2.3 Sellers' choices over time

The previous description showed that auctions are more likely to result in sales than posted prices and that sellers have different propensities to choose specific mechanisms even when conditioning on observable seller types. Further, it showed that both mechanisms are systematically employed, which can be a first indication that sellers do benefit from being able to choose from different selling mechanisms and especially auctions, since they correspond to more than two thirds of all listings (see Table 3).

However, this initial analysis ignored the dynamic component of the environment, namely whether and how sellers' choices varied as a function of distance to the deadline. This particular feature of seller behavior is relevant because the structural model is constructed around the sellers' dynamic optimization problem. This approach is appropriate provided that observed seller choices are consistent with forward-looking behavior. Hence, I now document the main empirical patterns associated with the evolution of three features of sellers' behavior over time: entry, mechanism and feature choices, and exit. Game day is normalized to be day 0 throughout this analysis and the model as well. These patterns are gathered in Figures 1 and 2.

Figure 1a displays the number of tickets that were listed for the first time on each day until the game as well as the fraction of these tickets that entered the market as auctions averaged across the 245 games in the sample. A few features become immediately apparent. First, most tickets enter the market closer to the game, even though there is relatively little entry activity on game day. Second, there seems to be a weekly pattern in entry, with the average number of entering tickets jumping every seven days before the game. This is possibly because the NFL follows a rigid weekly schedule, in which most games take place on Sundays, with usually one game on Thursday and one on Monday. Thus, it could be the case that at the end of each round sellers make entry decisions depending on the outcomes observed up to that point in time. Finally, the average fraction of auctions among entering listings follows a weakly increasing trajectory until ten days before the game, rising from 66.59% 28 days before the game to 76.12% just ten days before the game. It then starts decreasing, at first slowly (reaching 69.38% seven days before the game), accelerating within the week of the game (57.27% two days before the game) and dramatically vanishing on the day before the game and game day. This feature of the data at the very end of the time horizon is not surprising since auctions on eBay need to last for at least 24 hours. For completeness, Figure 1b displays the same quantities as in Figure 1a without restricting the listing events to the first time a set of tickets was listed. The patterns are overall very similar, which is expected given that, as shown in Table 4, most tickets are only listed once.

To further characterize seller relisting behavior, Figures 1c and 1d display, conditional on choosing to relist after a failed attempt to sell, which mechanisms sellers choose if their previous choice was an auction or a posted price, respectively. The new choices are displayed as a function of the day in which their previous attempt failed, that is, the day when the seller took down a posted price listing or the day when an unsuccessful auction ended. Hence, they break down the rows from Table 6 across the number of days until the deadline. In both cases, the patterns of relisting choices made far from game day suffer from a small number of observations, so focus should be given to relisting decisions made within three weeks of the game. Auctions are highly persistent with a weakly decreasing trend: the probability of continuing to use an auction decreases from 94.61% 21 days before the game to 83.94% four days before the game, when this probability dramatically decreases due to the aforementioned time constraints of using an auction. On the other hand, the probability of continuing to use posted prices does not present this monotone trend: it is first decreasing, going from 91.18% 19 days before the game to 64.29% six days before the game, when it starts rising again. While the time constraints that accompany auctions play a clear role in sellers' relisting behavior, the relative transition from posted prices to auctions starting around two weeks before the game cannot be explained by this feature alone.

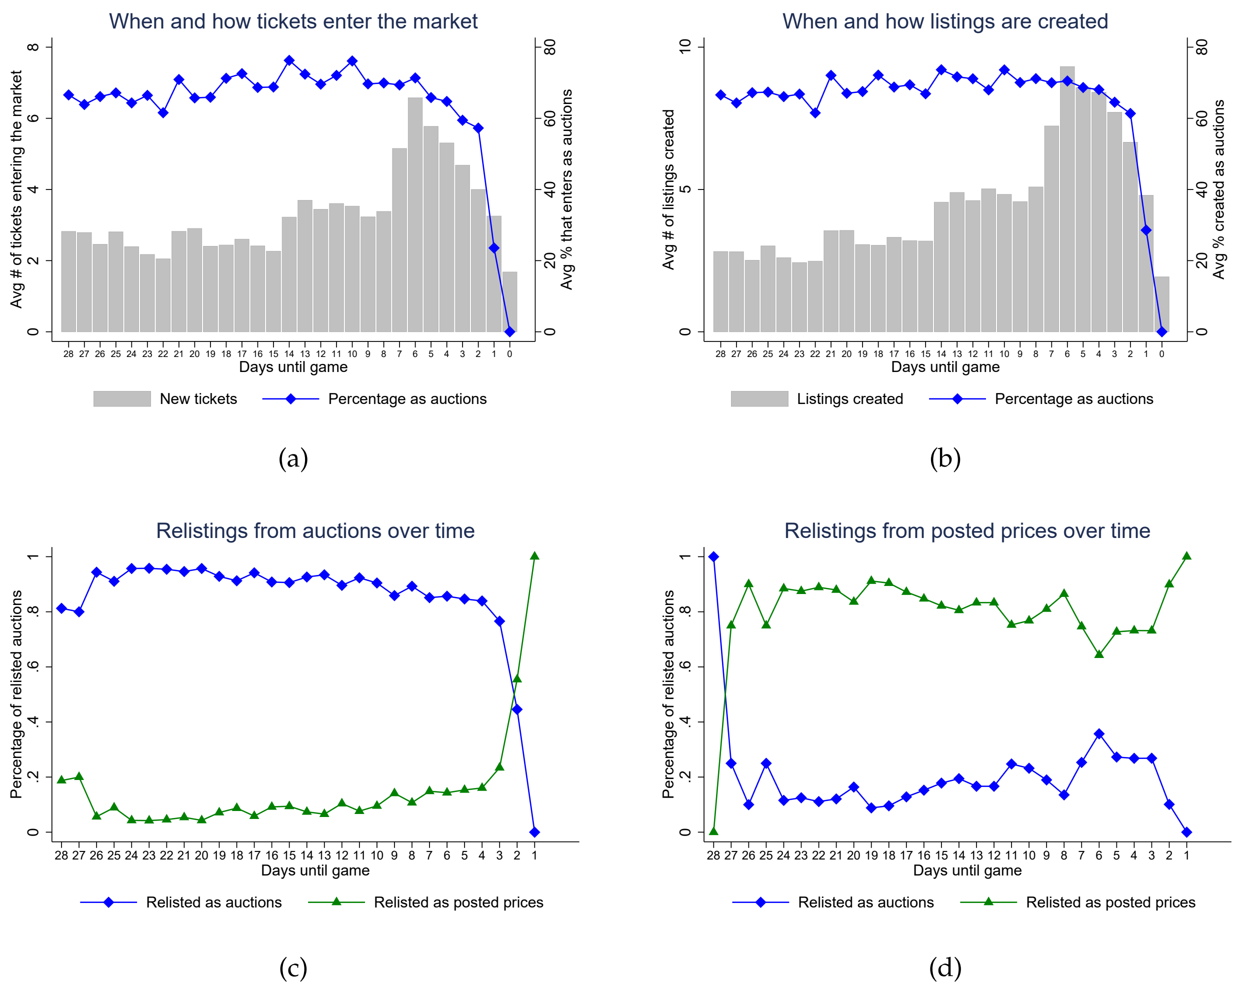

Figure 1: Features of supply evolution I

While Figures 1a and 1b provide measures of the flow of tickets into the website, a consumer who searches for listings is shown the entire stock of listed tickets at the moment of the query. Figure 2a displays the evolution of such stock over time, that is, the number of existing listings on each day until the game averaged across all 245 games in the sample, as well as the average fraction of existing tickets that are listed as auctions. The average stock increases day by day until five days before the game, which is a consequence not only of increased entry documented in Figure 1a but also of sellers switching to posted prices, which, unlike auctions, effectively have no binding deadlines. This switch can also be seen in Figure 2a through the average fraction of auctions, which is virtually constant until the week of the game, in part due to the aforementioned constraints on the use of auctions.

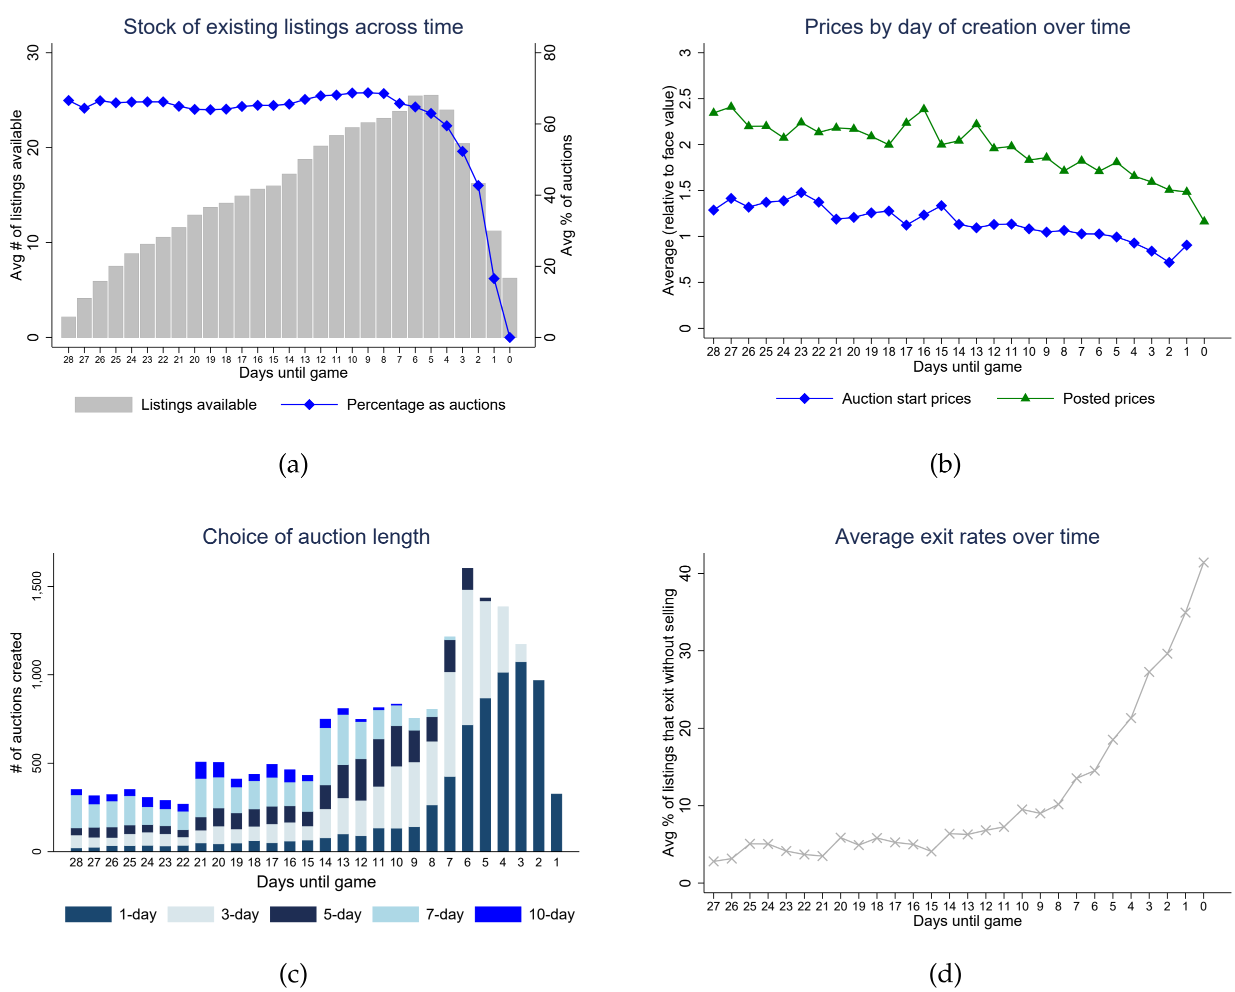

Figure 2: Features of supply evolution II

Having established when sellers enter the market and which mechanisms they choose to employ as a function of proximity to the deadline, I now document how the features of each mechanism are chosen. Figure 2b illustrates the evolution of prices. It plots the average posted and average reserve prices, with respect to the face value of the tickets, chosen for tickets listed on each day until the game. Both average reserve prices and average posted prices decrease as the deadline approaches, which is consistent with predictions from several revenue management models in which sellers are forward-looking.

The second feature sellers choose is the length of the auctions. Figure 2c shows the distribution of auction lengths among the auctions created on each day until the game summed over all 245 games in the sample. It is not surprising that once a format becomes infeasible sellers stop using it, that is, they do not create an auction that would have ended after the game had taken place. However, it is interesting to note that sellers start moving away from specific lengths even before they become infeasible. This suggests that when making a listing decision sellers do not only consider the current choice, but rather their entire selling strategy: choosing a longer length close to the deadline would imply a shorter period of time to try to sell their tickets again in case this attempt failed, which could explain the pattern observed in the data. This effect is reinforced if buyers also are forward-looking, because this would likely imply that they would become more hesitant to commit to longer auctions closer to the deadline, as not winning would imply they would have less time to try to acquire tickets from another source.

Having documented sellers' arrivals, mechanism choices, and mechanism feature choices, I now present the evolution of seller exit. To do this, I consider an exit opportunity to take place at the end of each day for a posted price listing that goes unsold and the expiration date of unsuccessful auctions. Figure 2d shows the evolution of the exit rate averaged across the 245 games in the data for each day until game day. Exit is weakly increasing until ten days before the game, rising from 2.8% 28 days before the game to 9% ten days before the game, and then increases at a much faster pace on the week of the game, peaking at 41.41% one day before the deadline. This pattern may be explained by a decreasing option value of waiting: the closer the game is, the lower the seller's expected payoff from selling in the future as the number of opportunities to do so diminish. If sellers have an outside option, they will choose to leave the platform if the expected value of remaining on it becomes sufficiently low. An alternative and reinforcing explanation is that outside options on game week potentially include actually going to the game. As demonstrated in the previously, the majority of sellers in the data are casual and therefore more likely to go to the game themselves, which could imply higher exit rates on game week.

So far the analysis has focused on the main empirical patterns regarding seller behavior solely as a function of the distance to the deadline. It could be the case that the evolution of such variables was driven by features other than proximity to the deadline, which would create a spurious relationship between them. To investigate whether this is the case I run several regressions controlling for ticket, game, and seller characteristics, average the residuals from these regressions for each day, and plot the evolution of the average residuals over time. Results are displayed in Figure 3 and preserve the patterns seen in the previous graphs. This indicates that such trajectories are not driven by spurious relationships between the variables of interest and the number of days until the deadline.

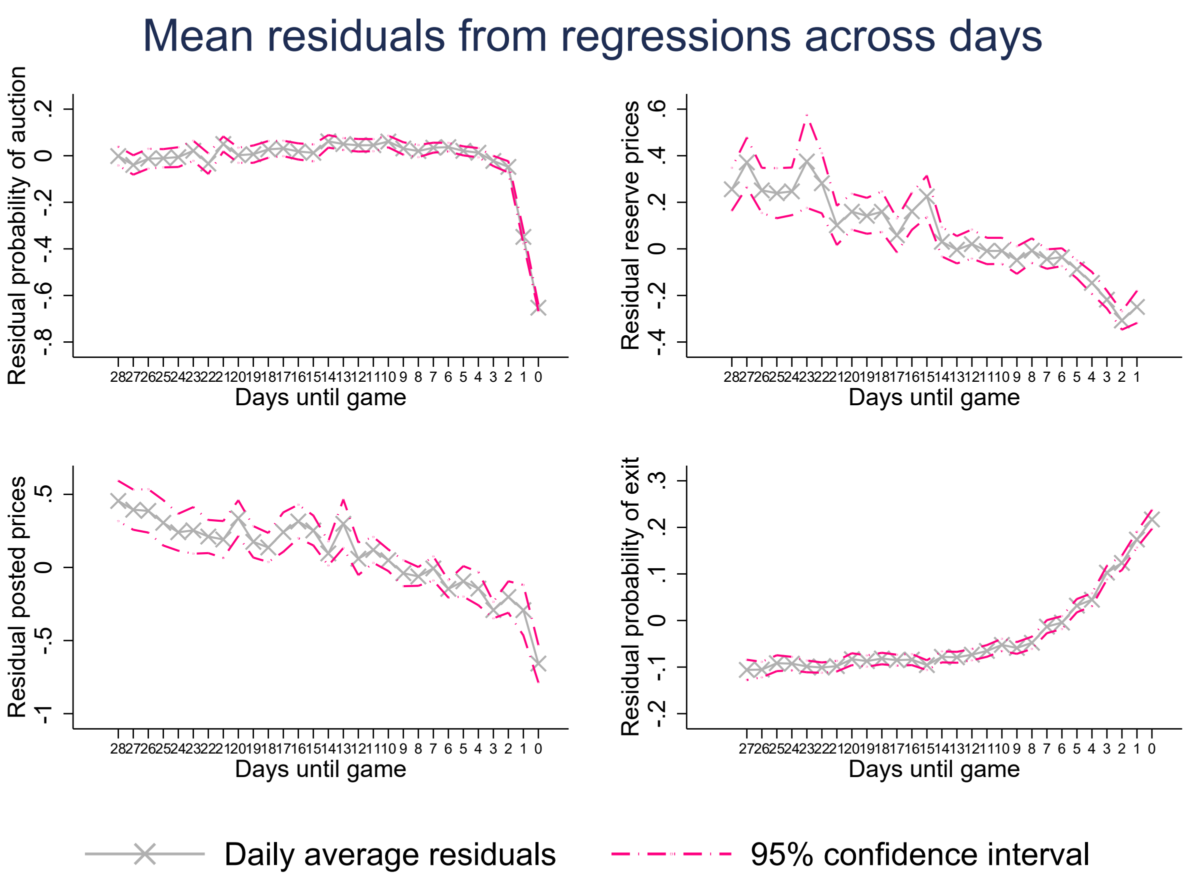

Notes: All regressions include game and ticket type fixed effects as well as indicators for the type of seller as defined in Section 2.2, the quartiles of seller's scores and ratings. Reserve and posted prices are expressed relative to the face value of the tickets. Standard errors used to construct the confidence intervals are clustered at the seller level.

Figure 3

Table 8: Effects of sellers' attributes on their choices

Regressor / Regressand	$\mathbb{1}\{auction\}$	Reserve price	Posted price	Auction length	Days at entry	$\mathbb{1}\{exit\}$	$\mathbb{1}\{sold\}$
Medium seller	0.03	0.08	0.11	-0.04	0.69	-0.01	-0.13
	(0.01)	(0.04)	(0.05)	(0.07)	(0.22)	(0.008)	(0.009)
Large seller	-0.006	-0.13	0.27	0.14	0.37	-0.03	-0.17
	(0.04)	(0.09)	(0.08)	(0.2)	(0.68)	(0.01)	(0.03)
2nd quartile of score	0.13	0.06	0.08	-0.92	-1.17	0.04	0.04
	(0.07)	(0.11)	(0.16)	(0.33)	(1.2)	(0.03)	(0.05)
3rd quartile of score	0.1	0.03	0.12	-0.71	-0.56	0.02	0.05
	(0.07)	(0.11)	(0.17)	(0.32)	(1.2)	(0.03)	(0.05)
4th quartile of score	-0.05	0.03	0.25	-0.87	0.97	-0.02	-0.03
	(0.08)	(0.12)	(0.17)	(0.33)	(1.24)	(0.03)	(0.06)
2nd quartile of rating	-0.03	-0.48	-0.44	0.74	1.22	-0.1	0.56
	(0.08)	(0.13)	(0.19)	(0.34)	(1.29)	(0.03)	(0.06)
3rd quartile of rating	-0.06	-0.49	-0.57	0.81	2.11	-0.09	0.56
	(0.08)	(0.12)	(0.18)	(0.34)	(1.27)	(0.03)	(0.06)
4th quartile of rating	-0.04	-0.49	-0.51	0.66	2.47	-0.09	0.59
	(0.07)	(0.11)	(0.17)	(0.32)	(1.19)	(0.03)	(0.05)
Mean dependent variable	0.68	1.08	1.86	3.56	11.74	0.15	0.55
Observations	28,257	19,082	9,175	19,082	19,174	48,547	19,174
Clusters	8,081	5,797	3,715	5,797	8,081	6,343	8,081
R-squared	0.1	0.29	0.43	0.49	0.07	0.13	0.29

Notes: All regressions also include game and ticket types fixed effects. All but the fifth also include indicators for the day when the action chosen by the seller took place. Auction start prices and posted prices are calculated with respect to the face value of the tickets. Standard errors, displayed below the coefficients between parentheses, are always clustered at the seller level.

A final relevant matter is whether different types of seller make systematically different choices. To investigate this I run several regressions of decision variables by sellers on indicators of their type as defined in Section 2.2 and on the quartiles of sellers' score and rating, with coefficients reported in Table 8. Even though there does not seem to be a consistent and systematic pattern, a few relationships become apparent. Medium sellers are more likely to arrive to the market earlier, to use auctions, and to choose relatively higher reserve prices. Medium and large sellers also choose relatively higher posted prices and are less likely to sell. Large sellers are also less likely to leave the market without selling. It can also be seen that sellers with higher feedback scores tend to choose shorter auctions, while sellers with higher ratings charge lower reserve and posted prices, pick longer auctions, are less likely to leave without selling, and more likely to sell. Overall these results suggest that observable seller heterogeneity is relevant in explaining their decisions, and therefore will be accounted for in estimating the structural model.

2.4 Buyer arrival, exit, and clicking behavior

The previous description showed how sellers' decisions change as a function of distance to the deadline. The empirical patterns show that sellers' choices are consistent with forward-looking behavior, which supports the modeling approach of framing their decisions as arising from a dynamic mechanism choice problem. Further, it showed that their choices were systematically correlated with observable seller attributes such as their type and feedback scores, which informs subsequent model specification. I now perform a similar analysis, but focusing on the demand side of the market instead. Both analyses in part intend to simply describe how agents behave in a market for perishable goods with mechanism choice. This can be valuable on its own, and especially on the demand side because of the novel availability of clickstream data, which provide a better measure of potential demand than what the majority of previous studies had.

Nevertheless, there also is a difference between the objectives of these two analyses. One of the goals of the previous analysis was to show that the observed sellers' choices were consistent with forward-looking behavior to support the modeling decision of framing the sellers' problem as a dynamic mechanism choice one. However, in line with standard revenue management models, the empirical model presented below treats demand in a highly stylized fashion for the sake of tractability. This treatment is in all likelihood inconsistent with many of the empirical patterns shown here. It is my goal to be transparent about this tension.

Following buyers on an online platform without access to browsing data is a harder task than keeping track of sellers or listings. I only observe a potential buyer on the website when this buyer clicks to view a listing's detail page. This is not a perfect measure because all buyers who arrive to the platform, type a query for game tickets, observe the results but choose not to click on any of them are ignored. However, it does convey information about the activity level on the demand side of the market, which I now summarize and display in Figure 4.

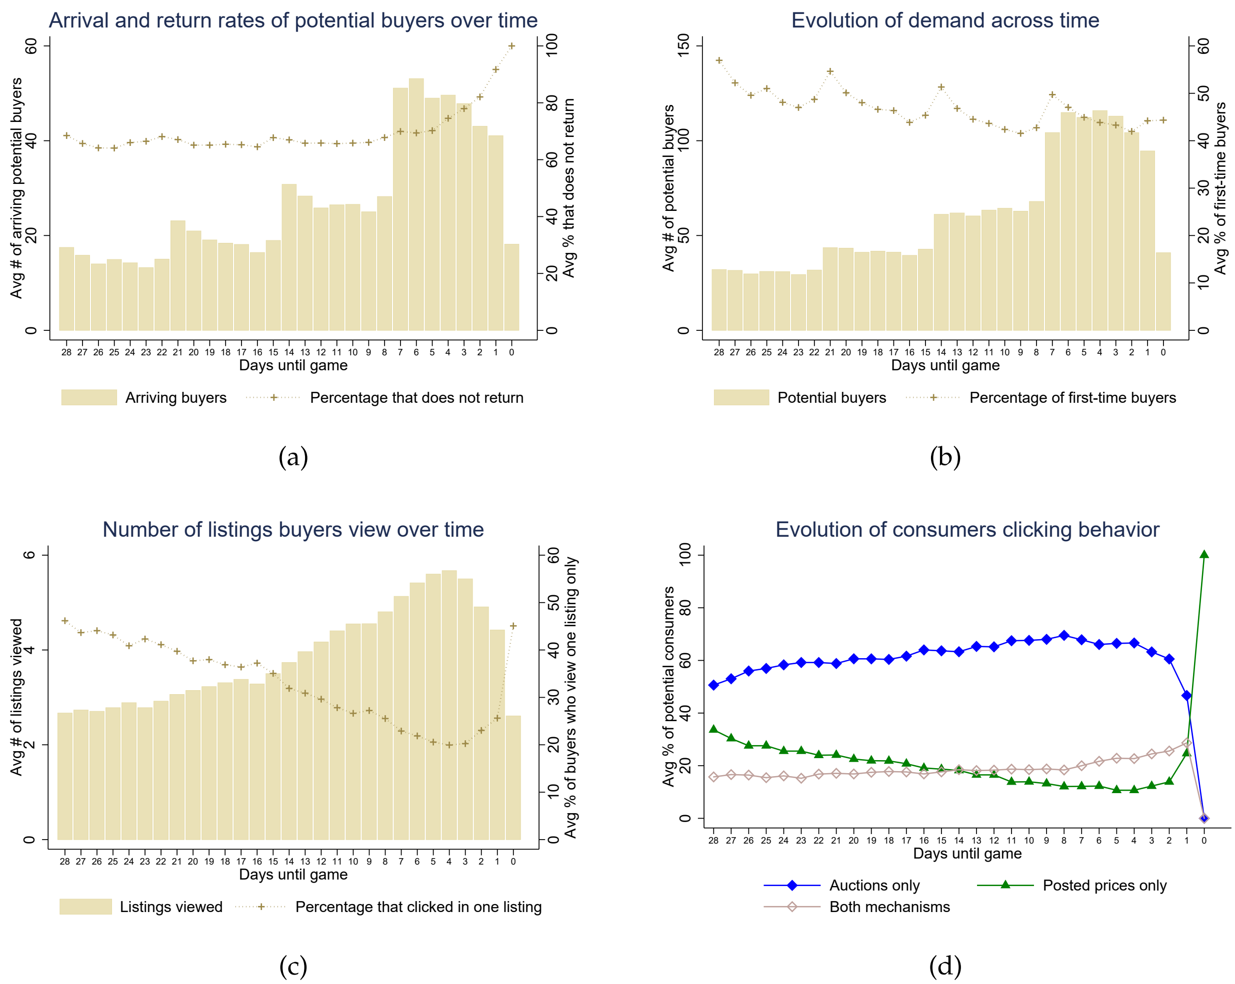

(a) (b) (c) (d)

Figure 4: Features of demand evolution

Figure 4a displays the number of potential buyers that enter the market on each day, that is, that are observed clicking to view a listing for the first time on each day. This quantity is averaged across all 245 games in the sample and features the same weekly pattern seen in Figure 1a, suggesting that both buyers and sellers make decisions depending on how the season progresses. In addition, Figure 4a displays the average fraction of entering potential buyers that do not return in the future. This fraction is high and stable until 5 days before the game, ranging between 65% and 70%, when it starts increasing quickly. In addition, on average just around 4.2% of these buyers leave after making a purchase, which suggests that most buyers on this market might actually be short-lived, arriving and leaving on the same day. To address the overall size of demand and the fraction for which entering buyers are responsible, Figure 4b displays the average number of all buyers in the market seen on each day and the fraction of them who was seen in the market for the first time. The weekly pattern once again becomes apparent as well as the fact that the market becomes more active as the deadline approaches. However, it also shows that roughly half of buyers on the market are returning ones, which can be consistent with forward-looking behavior.

Having illustrated when buyers arrive and leave, I now document their clicking behavior in more detail. Figure 4c shows the average number of listings buyers view on each day. This average increases from around three listings three weeks before the game to almost six just four days before the game, when it starts decreasing. The increase in the number of listings viewed is partly a consequence of the increase in the overall supply of tickets as displayed in Figure 2a. It is also interesting to note the decrease in the average fraction of buyers who only click on one listing: it goes from more than 46% a month before the game to less than 20% four days before the game, when it starts increasing until the deadline is reached. This latter trend can be a consequence both of the overall lower supply of tickets (Figure 2a) and the shorter amount of time individuals have to potentially acquire tickets.

Finally, Figure 4d displays what buyers click on. The average fraction of potential buyers who only click on auctions increases from a little less than 50% a month before the game to almost 70% just eight days before the game. This is an interesting phenomenon given that, as shown in Figure 2a, the average fraction of auctions among existing listings is virtually constant during this timespan. Symmetrically, the fraction of potential buyers who only view posted prices is decreasing until four days before the game, while the fraction of buyers who view both mechanisms is relatively stable until the week of the game, ranging between 15% and 19%, when it the rises up to almost 29% and then vanishes on game day due to the time constraints associated with auctions.

2.5 Market evolution and outcomes

The previous two analyses described how the supply side and the demand side of the market evolved separately. For completeness, I now describe how the market as one changes as the deadline approaches, presenting such patterns in Figure 5.

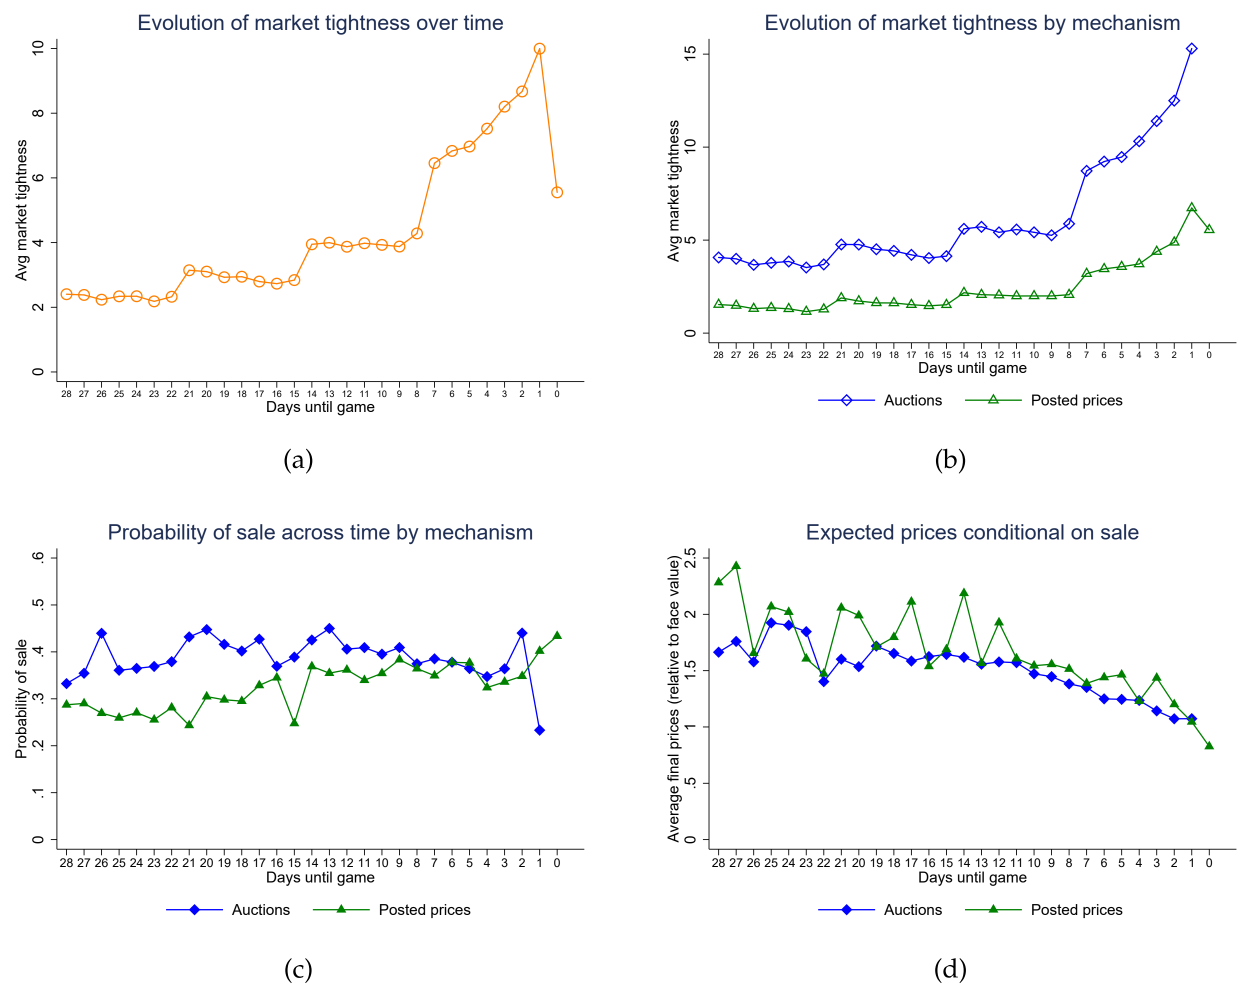

Figure 5: Features of overall market evolution

Figure 5a displays the evolution of market tightness over time. Market tightness is defined as the ratio between the number of buyers on the market and the number of existing listings available, and it will play a key role in the model to explain how buyers are matched to listings. The aforementioned weekly pattern can be seen again, and the fact that average market tightness is increasing (rising from a little more than 2 four weeks before the game to more than 10 the day before the game) suggests that market conditions become more favorable to sellers as the deadline approaches, as it indicates that the number of potential buyers for each existing listing is higher. In other words, it indicates that demand grows faster than supply in this market.

Given that most buyers seem to join either the auction market or the posted price market, as suggested by Figure 4d, I now separate the evolution of market tightness into these two mechanism markets and present the results in Figure 5b. The patterns by mechanism are qualitatively similar to the overall market tightness, even though the auction market not only has higher average tightness levels but also experiences a more dramatic increase as the deadline approaches. In theory, this indicates that the auctions become more advantageous closer to the deadline, as their ability to aggregate information, discover prices, and to extract consumer surplus increases with the number of buyers.

Finally, I describe the performance of each mechanism. From the vantage point of a seller, a mechanism can roughly be evaluated using two dimensions: the probability of a sale and the expected price conditional on a sale. Figures 5c and 5d display the evolution of the probability of sale and expected transaction price given a sale over time across mechanisms, respectively. The probability of sales is reasonably stable over time for both auctions and posted prices despite the increase in market tightness seen in Figure 5b. On the other hand, average transaction prices decrease over time for both auctions and posted price listings, which should be expected given that listed reserve prices and listed posted prices decrease over time, as displayed in Figure 2b. However, a notable feature of the data is that auctions are more likely to sell while posted prices yield higher revenues conditional on a sale, and that this dichotomy effectively holds at every point in time. This implied tradeoff between auctions and posted prices is not an exclusive feature of these data as it was also found in several other contexts.[3]

3 Empirical model

I now present the empirical model that is estimated using the data described above. To perform counterfactual simulations in which existing selling mechanisms from the sellers' choice sets are removed, this model aims primarily at rationalizing sellers' mechanism choices. I begin by describing the demand side of the model, going over the technology that matches buyers to listings and buyer's expected payoff from a posted price and an

[3]Examples include Hammond (2010), Sweeting (2013), Bauner (2015), Einav et al. (2018), and Coey et al. (2020).

auction listing. I then proceed to describe the supply side, and conclude by defining what a market equilibrium is in order to close the model.

3.1 Demand

Since the focus of this paper is on sellers' dynamic behavior, the demand side of the model is relatively more stylized. When all assumptions and modeling choices are put together, they effectively imply that demand follows an exogenous statistical process, which is in line with the approach taken by many revenue management models. I chose this approach solely to keep the model simple and tractable as it will be detailed below.

Buyers are assumed to be risk-neutral, short-lived and to exogenously arrive to the market. Upon arrival, a buyer's only decision is whether to join the auction market or the posted price market. Conditional on joining one of these markets, a buyer is matched with a single listing and his decision is to either accept the posted price if in the posted price market or to submit a bid if in the auction market. More explicitly, I maintain the following assumption.

Assumption 1.
(i) Buyers are risk-neutral, randomly arrive to the platform and leave the market after their purchase opportunity ends whether they make a purchase or not. Buyer i's willingness-to-pay for listing j in order to acquire it on day t, V_{ijt}, is a private and independent draw from a distribution $F_V(\cdot|X_j, t)$, in which X_j is a vector of listing j's characteristics.
(ii) Each buyer is matched to only one listing after they choose which market to join. Valuations are drawn after buyers are matched.

Assumption 1 is made solely to keep the model tractable and warrants several comments, beginning with part 1(*i*). Given the patterns seen in Figure 4, it is straightforward to see that this assumption is not supported by the data: while many buyers are indeed only seen in the market on one day, many are observed clicking on many different days. Assumption 1(*i*) can be interpreted as implying that activities by the same buyer across different days are independent from one another, which is, in all likelihood, implausible, not to mention that many buyers are only seen once in the platform because they choose not to return. While many studies have allowed for forward-looking buyers in the context of dynamic auctions, such as Hendricks and Sorensen (2018), Backus and Lewis (2020), and Bodoh-Creed et al. (2021), and in the context of durable goods, such as Nair

(2007), Goettler and Gordon (2011), and Gowrisankaran and Rysman (2012), they have all relied on stationarity conditions to maintain tractability. However, my setting is inherently nonstationary given that it is a market for perishable goods with an exogenous and commonly shared deadline, as well as with the possibility of bidding. Allowing for forward-looking buyers under these circumstances would possibly render the estimation of this model infeasible. To partially capture forward-looking buyer behavior, I allow the distribution of buyer valuations to exogenously change over time. In estimation, I specify that the temporal changes in the distribution valuations are not interacted with the vector of ticket characteristics, X_j. Although parsimonious, this specification does impose restrictions on the model. First, in a model with forward-looking buyers and a deadline, time itself becomes a state variable, and therefore should be interacted with the vector X_j, which is why this specification only captures forward-looking behavior partially. Second, it rules out any uncertainty over buyer valuations over time, which could arise from realizations of game-specific random events. These events could be, for example, injuries sustained by important players or games becoming more or less relevant for the purposes of qualifications to playoffs, which is a function of different game results.

Part 1(*ii*) is also challenged by the patterns seen in Figure 4. First, the fact that most buyers click on several listings challenges the idea of them being matched to a single listing. Second, a substantial fraction of buyers click on both auctions and posted prices, which further challenges the assumption that they choose to enter one mechanism market. Moreover, assuming that valuations are only drawn after a match is made precludes buyers from self-selecting into mechanism markets based on their valuations, which is a strong condition. Nevertheless, all these pieces play key roles in the identification and estimation strategies I follow.

Under Assumption 1, upon arrival a buyer needs to decide, before observing his valuation, which market to join, auctions or posted prices. To make this decision, the buyer, who maximizes his expected payoff, needs to compute his expected utility from joining each of the markets. In turn, to compute these expectations buyers need to have beliefs over posted prices (P), reserve prices (R), ticket characteristics on each market (X), and auction end dates (T), because they can be matched to an auction that ends days later and because, under Assumption 1(*i*), their willingness-to-pay is determined by the day in which they would acquire the tickets, and not their date of arrival to the market. In addition, buyers need to have beliefs about their competition, that is, the arrival rates of other buyers to listings, which I denote by Λ. Finally, buyers need to know the distribution from which valuations are drawn. The assumption below states that buyers' beliefs correspond to the true probability distributions of these objects.

Assumption 2.
After arriving to the market on day t, buyers know: the joint distributions of posted prices, arrival rates, and characteristics of listings available in the posted price market, $F^{P_t}_{P,\Lambda,X}(\cdot)$; the joint distribution of reserve prices, end dates, arrival rates, and characteristics of listings available in the auction market, $F^{A_t}_{R,T,\Lambda,X}(\cdot)$; and the distribution from which valuations are drawn, $F_V(\cdot|X,t)$.

I now explain in detail how buyers make their decision of which market to join. First, I describe what is the matching technology that implies the arrival rates indexed by Λ. I then consider the expected utility buyers accrue from participating in auctions and posted prices. Finally, I explicitly state what their optimal choices are.

3.1.1 Matching and market operation

The marketplace consists of two separate sub-markets defined by the mechanism used in them, auctions or posted prices. These markets are indexed by k, where $k \in \{A, P\}$. I assume that the number of potential buyers attracted by listing j in market k at day t is drawn from a Poisson distribution with parameter Λ^k_{jt} and that these draws are independent over time. The Poisson arrival rate assumption, which is common in revenue management models, is maintained here solely for the tractability it gives both to the supply and demand sides of the model, as will become clear below. When compounded with the independent draws over time assumption, it particularly facilitates the analysis of auctions with different lengths because it implies that if listing j is an auction of length ℓ created on day t, its total arrival rate is given by $\Lambda^{A_\ell}_{jt} \equiv \sum_{d=0}^{\ell-1} \Lambda^{A}_{j,t-d}$. This is important because, as it will be made clear below, this model endogenizes the auction length choice, unlike other similar models in the mechanism choice literature or that study seller behavior on eBay.

To motivate the Poisson assumption and inform its specific parametrization, I resort to an asymptotic approximation in a directed matching setting.[4] Let S^k_{jt} be the overall number of existing listings of tickets for the game to which listing j is associated in market k on day t, which is determined by sellers' choices as explained below. Further, let B^k_{jt} be the total number of buyers in market k for the game to which a listing j is associated on day t. I assume that a buyer is matched with listing j with probability $\frac{\lambda^k_{jt}}{S^k_{jt}}$, where λ^k_{jt} is

[4]For an introduction to this topic, see, for example, Rogerson et al. (2005).

smaller than S^k_{jt}. If the parameter λ^k_{jt} is equal to one for all j, this assumption would imply that a buyer is matched to one of the sellers in market k on day t at random. I allow λ^k_{jt} to vary across sellers to capture observable seller characteristics that explain why some sellers might attract more or less buyers.

Under this matching scheme, the probability that listing j in market k on day t does not attract any buyers is given by $\left(1 - \frac{\lambda^k_{jt}}{S^k_{jt}}\right)^{B^k_{jt}}$. I now consider what happens to this expression in a large market with many buyers and sellers: taking the limit of this expression as both B^k_{jt} and S^k_{jt} diverge to infinity and assuming that $\frac{B^k_{jt}}{S^k_{jt}} \to \alpha^k_{jt}$, then $\left(1 - \frac{\lambda^k_{jt}}{S^k_{jt}}\right)^{B^k_{jt}} \to e^{-\lambda^k_{jt}\alpha^k_{jt}}$. This limit is consistent with a Poisson specification because the probability of attracting no buyers, $e^{-\lambda^k_{jt}\alpha^k_{jt}}$, is equal to the probability of a Poisson random variable with parameter $\lambda^k_{jt}\alpha^k_{jt}$ taking the value zero. Thus, I specify $\Lambda^k_{jt} = \lambda^k_{jt}\alpha^k_{jt}$. In this specification, α^k_{jt} is the market tightness of market k on day t for the game to which the tickets in listing j are associated, as displayed in Figures 5a and 5b.

This approach to matching is the most simplified part of the model. It leverages Assumption 1(*ii*), which, as discussed above, is not supported by the data. In addition, it virtually rules out any active buyer search from this model. Nevertheless, this approach is taken due to the amount of tractability it provides. It summarizes supply and demand conditions into a single variable, which, as it will be made clear below, becomes the only state variable along with time in the sellers' dynamic mechanism choice problem.

Having outlined how buyers are matched to listings, I now derive a buyer's expected payoff from an auction and from a posted price. Under Assumption 2, conditional on a match buyers have all the information required for them to compute their expected payoffs. Here is where one the main benefits of the Poisson arrival rates becomes apparent, namely the "environmental equivalence" demonstrated by Myerson (1998): if the overall arrival rates of buyers to a listing is Λ, it is clear that buyers and sellers will use Λ directly to compute their expected profits. However, environmental equivalence implies that when a given buyer is assessing his expected payoffs and therefore his competition, the same arrival rate Λ defines the statistical process according to which rival buyers will arrive to the listing to which this buyer is matched. This simplifies matters because the same arrival rate will therefore be used by sellers and buyers in their respective decision-making problems.

3.1.2 Buyer's expected utility from auctions

Before deriving a buyer's expected utility from auctions, I first state more explicitly how I assume that auctions are conducted within the model.

Assumption 3.
At the end date of an auction, all buyers who were matched to it are randomly ordered and called to submit their bids one by one. At the time buyer i has to submit his bid, the only information he has about competition is the highest losing bid, which is equal to the reserve price until two bids above it are submitted. The winner is the buyer who submits the highest bid, and he pays the highest losing bid.

Assumption 3, when compounded with Assumption 1, which defines this model as one where buyers have symmetric independent private values (IPV), implies that it is optimal for a buyer to bid his own valuation unless the highest losing bid exceeds it, in which case no bid is submitted. Akin to Song (2004), Assumption 3 implies that the auctions in this model are outcome equivalent to a sealed-bid second-price auction while allowing for a subset of bidders not to submit any bids, which is a feature of online auctions. However, it is important to note that, as outlined by Zeithammer and Adams (2010), this approximation is not inconsequential and does not correspond to how e-commerce auctions, such as those conducted on eBay, occur in reality. Assumption 3 is nevertheless maintained to keep the model tractable and because of its consequences for the identification and estimation strategies I adopt, which will be outlined below.

I now present a buyer's utility from participating in an auction. Assume that buyer i arrives at day t, with valuation v, and is matched with an auction that has reserve price r, end date τ, characteristics x, and Poisson arrival rate λ. Because of environmental equivalence, i's expected utility, $U_{A_t}(v, r, \tau, \lambda, x)$, is given by (see Appendix C for the derivation):

$$U_{A_t}(v, r, \tau, \lambda, x) = \int_r^v e^{-\lambda[1-F_V(u|x,\tau)]} du \tag{1}$$

if $v > r$ and zero otherwise. Buyer i has to integrate this expression with respect to all its features, which is enabled by Assumption 2. For the purposes of integrating it with respect to v, this requires i to integrate over the features of the listings available (x) on the auction market and their end dates (τ). Thus, the final expected utility is given by:

$$\bar{U}_{A_t} = \int_{R,T,\Lambda,X} \left(\int_r^\infty U_{A_t}(v, r, \tau, \lambda, x) dF_V(v|x, \tau) \right) dF^{A_t}_{R,T,\Lambda,X}(r, \tau, \lambda, x). \tag{2}$$

3.1.3 Buyer's expected utility from posted prices

Now assume that i is matched to a posted price listing on day t with price p, arrival rate λ, and characteristics x. Notice that the end date is no longer relevant since a posted price listing is evaluated daily. Hence, I maintain the following assumption about how posted price listings work in this market that is analogous to Assumption 3.

Assumption 4.
At the end of each day, all buyers who were matched to a posted price listing are randomly ordered and called to either accept or reject the posted price.

Environmental equivalence still holds, so buyer i's expected utility, $U_{P_t}(v, p, \lambda, x)$, assuming his valuation is above the posted price, is given by (see Appendix C)

$$U_{P_t}(v, p, \lambda, x) = \frac{(v-p)\left(1 - e^{-\lambda[1-F_V(p|x,t)]}\right)}{\lambda\left[1 - F_V(p|x,t)\right]}. \tag{3}$$

If the buyer's valuation is below the posted price, his utility is simply zero. Once again, buyer i has to integrate such expression with respect to i's own valuation, v, which requires integration with respect to the characteristics of the listing, x, as well as with respect to prices and arrival rates. Thus, the final expected utility is

$$\bar{U}_{P_t} = \int_{P,\Lambda,X}\left(\int_p^{\infty} U_{P_t}(v, p, \lambda, x) dF_V(v|x,t)\right) dF^{P_t}_{P,\Lambda,X}(p, \lambda, x). \tag{4}$$

3.1.4 Buyer's market choice

Given the assumptions and objects defined above, the buyer's decision at the time of his arrival is straightforward. Buyer i's utility from joining market k when arrives on day t is

$$U_{ikt} = \bar{U}_{k_t} + \eta_{ikt},$$

where η_{ikt} is a random shock with full support and independent across buyers, markets, and time. Hence, a buyer will join market k instead of market k' if and only if

$$U_{ikt} \geq U_{ik't}, \tag{5}$$

which happens with probability $F_{\Delta_\eta}\left(\bar{U}_{k_t} - \bar{U}_{k'_t}\right)$, where $F_{\Delta_\eta}(\cdot)$ is the cumulative distribution function of $\eta_{ik't} - \eta_{ikt}$.

3.2 Supply

Having outlined the demand side of the model, I now proceed to the supply side. Before stating sellers' choices mathematically, I first describe the main elements of their dynamic mechanism choice problem informally.

Forward-looking sellers are assumed to be risk-neutral and to arrive to the market at random, so that seller entry is exogenous. Upon arrival, sellers know precisely how all payoff-relevant variables will change over time. Having all the relevant information to compute her expected payoffs, a seller chooses between listing her tickets as an auction or a posted price. The seller can choose auctions of five different durations, which correspond to the options available at eBay: one, three, five, seven, or ten days long. In addition, the seller needs to decide the reserve price for the auction. In turn, the only feature the seller chooses when listing her tickets as a posted price is the posted price itself.

Under common circumstances, an expected payoff maximizing seller would always choose an auction such as the one in this model over a posted price. This is because, all else constant, for any posted price p, the seller could instead run an auction with a reserve price equal to p, sell with equal probability and earn a weakly higher expected payment conditional on a sale due to the competitive bidding induced by the auction. However, the dynamics of this setting make it possible for this result not to hold. When listing her tickets, a seller has to pay a listing, or convenience, cost. Whenever the listing goes unsold, the seller needs to consider whether to relist her tickets and pay the listing cost again, or exit the market and earn her outside option. One of the benefits of creating a posted price listing is that, as long as the seller does not alter the posted price itself, the listing does not expire and the seller does not have to pay the listing cost to keep her tickets available on the platform. In turn, auctions always have an ending date, and to relist the tickets after an unsuccessful auction the seller has to pay the listing cost. Therefore, all else constant a seller with a higher listing cost is more likely to choose to list her tickets as a posted price instead of as an auction due to this convenience. This is also why exploiting differences in choices of auction length is valuable, as they can also be informative of the magnitude of a seller's listing costs.

To formalize this model, I begin by stating the following assumption regarding seller preferences, entry, and supply, which is maintained throughout the paper.

Assumption 5.
Sellers are risk-neutral, enter the market randomly, and have one pair of tickets to sell by day 0 (game day). There is no discounting.

Seller risk neutrality is assumed solely for the sake of simplicity, as is the lack of discounting. However, the latter can be supported by the fact that this model addresses a daily decision-making problem over a relatively short time horizon (one month) in which there is not a stream of realized payoffs, but rather a one shot payment that is realized whenever a sale is made. In turn, the assumption that each seller has only one pair of tickets to sell allows me to use the subscript j to denote both sellers and tickets interchangeably. In the data almost 60% of sellers in the final sample indeed offer only one set of tickets per game. Whenever the same seller holds multiple pairs of tickets, this is equivalent to assuming that the seller's decisions regarding each pair are independent, ruling out inventory management considerations that would add substantial complexity to the model. Finally, the assumption that seller entry is random is not supported by the data as the fifth column of Table 8 indicates that seller attributes correlate with their arrival date. This will in part be captured by allowing seller parameters to vary with seller characteristics in estimation. Nevertheless, this condition becomes more restrictive in the context of the counterfactuals exercise of interest, in which mechanisms are eliminated from the platform, more specifically auctions. As mentioned before, eBay is the main platform for tickets where auctions are available, so it could be expected that sellers (and buyers) come to eBay precisely due to the availability of using auctions. The assumption is kept for the sake of simplicity, but this caveat should be kept in mind and will be brought up again in the discussion of counterfactual results.

Sellers' observable characteristics are gathered in two vectors, X_j and Z_j, which are, respectively, associated with ticket characteristics such as their location within a stadium and attributes that attract potential buyers such as feedback rating. They can choose between two mechanisms, auctions, denoted by A_ℓ, where ℓ denotes the auction's duration, and posted prices, denoted by P. As above, mechanisms are also denoted with k, so that $k \in \{(A_\ell)_{\ell \in \mathbb{L}}, P\}$, where $\mathbb{L}$ is an exogenous and finite set from which sellers can choose the auction duration.[5] Each k entails a choice of price, which is the reserve price for auctions or simply the posted price. In addition, sellers need to pay a seller-specific listing

[5]This set is the same as the one offered by eBay, so that $\mathbb{L} = \{1, 3, 5, 7, 10\}$.

cost, κ_j, each time they list their tickets. Hence, it can also be interpreted as a monitoring cost or an inconvenience disutility. Accordingly, if a seller creates a posted price listing but does not alter the posted price itself, this cost need not be paid again after the tickets are first listed.

I now proceed to describe a seller's payoff from each mechanism, followed by a characterization of her optimal choice. From each individual seller's perspective, given the matching process described above, the number of potential buyers who randomly arrive to j's listing on day t when the chosen mechanism is k is N_{jt}^k. This is a Poisson random variable with the aforementioned arrival rates of $\Lambda_{jt}^k = \lambda_{jt}^k \alpha_{jt}^k$. The arrival rates depend on characteristics of the seller, Z_j. In particular, I specify that $\lambda_{jt}^k = \exp\{Z_j' \lambda_t^k\}$, which are not necessarily the same as the ones that affect bidders' valuations above. Furthermore, they depend on the market tightness level at day t for each mechanism k, α_{jt}^k. For a forward-looking seller to make decisions, therefore, she needs to have beliefs over how market conditions, namely market tightness on each mechanism market, evolve over time. The following assumption states that sellers know perfectly how such market conditions evolve.

Assumption 6.
Sellers have perfect foresight over the distribution of buyer valuations and over market tightness in both markets.

Assumption 6 requires sellers not only to know how buyers' distribution of valuations change over time, but also to know how market tightness evolves in each market. Consider a seller assessing the expected payoff from holding a 5-day auction. This seller needs to have expectations regarding how many potential buyers will arrive during these five days. If future market conditions were unknown, this seller would need to have expectations regarding not only the state of the market on the next day, but on the next four days as well, which would make the seller's optimization problem considerably more complex. Nevertheless, this requirement might be more stringent when it comes to smaller, casual sellers, who arguably are less experienced and not as familiar with the market. However, the fifth column of Table 8 indicates that such sellers arrive closer to the deadline, which partially mitigates this concern. Hence, I maintain this assumption solely for simplicity.

A different concern is the choice to abstract from competition between sellers, effectively treating their optimization problem as a single-agent one. While directly modeling strategic interactions between sellers would bring more richness to the model, as noted by Sweeting (2013) modeling a perishable good market with mechanism choice as a game would almost surely make its estimation infeasible. Other studies that attempted to model mechanism choice such as Hammond (2013), Sweeting (2013) and Bauner (2015) treated the seller's problem as a static one and only introduced competition in a reduced-form way, while Coey et al. (2020) pose a model with a continuum of sellers that earn zero profits in a steady-state equilibrium. Lee et al. (2012), Sweeting and Sweeney (2015) and Sweeting (2015) did estimate dynamic models of perishable good markets, but only focused on pricing.

I now present the mechanism-specific value functions and sellers' final choices.

3.2.1 Auctions

Before deriving a seller's expected payoff from auctions, I first state more explicitly what I assume regarding sellers' commitment power.

Assumption 7.
If seller j chooses an auction, the seller stays locked into this auction until it ends.

In the data sellers very rarely end auctions prematurely. While in practice sellers can end auctions early without having to sell even if a buyer has already submitted a bid above the reserve price, as can buyers retract already submitted bids, these decisions yield reputational and sometimes monetary costs, which is arguably why they are rarely observed in the data. Assumption 7 is analogous to Assumption 3 in that they imply that buyers and sellers become "locked into" auctions, which, in turn, allows me to treat them as if they were sealed-bid auctions, thereby ruling out intra-auction dynamics. However, it is important to note that this does not entirely correspond to how eBay auctions are conducted in practice. I place these restrictions on the model solely for tractability purposes.

Given the constraints on how auctions work, a seller's expected payoff from creating an auction of length ℓ on day t is as follows. The number of buyers who will be matched to j's listing is given by $N_{jt}^{A_\ell} = \sum_{d=0}^{\ell-1} N_{j,t-d}^{A}$. This is a random number from the perspective of the seller, who will, therefore, have to integrate with respect to it. Since each $N_{j,t-d}^{A}$ follows a Poisson distribution and each draw is independent from the others, it follows that $N_{jt}^{A_\ell}$ also follows a Poisson distribution, with parameter $\Lambda_{jt}^{A_\ell} = \sum_{d=0}^{\ell-1} \Lambda_{j,t-d}^{A}$. Notice that I assume that different auction lengths do not attract potential buyers differently in any other way rather than the duration itself. That is, all auctions on a given day attract buyers according to the same distribution. Omitting the conditioning variables to ease notation, the value function for such an auction is given by

$$\pi_{jt}^{A_\ell} = \max_r \mathbb{E}_{N_{jt}^{A_\ell}} \left[\Pr\left(V_{j,t-\ell+1}^{(n:n)} \geq r\right) \mathbb{E}\left[\max\{V_{j,t-\ell+1}^{(n-1:n)}, r\} \middle| V_{j,t-\ell+1}^{(n:n)} > r\right] \right.$$
$$\left. + \Pr\left(V_{j,t-\ell+1}^{(n:n)} < r\right) \Pi_{j,t-\ell} \middle| N_{jt}^{A_\ell} = n \right] - \kappa_j + \epsilon_{jt}^{A_\ell} \equiv \tilde{\pi}_{jt}^{A_\ell} + \epsilon_{jt}^{A_\ell}. \tag{6}$$

The first term in (6) simply states that a purchase will be made if the highest valuation among buyers that arrive to j's listing is above the reserve price chosen by j, where the superscript $(n : n)$ indicates, following the usual notation from order statistics, the largest value out of n values. Since the model implies that bidders will play their weakly dominant strategy and bid their valuations, the expected payoff conditional on a sale is simply the greater of the second highest valuation and the reserve price. Sellers have to pay a listing cost of κ_j at the moment the listing is created. In case the item is not sold, the seller has a continuation value, $\Pi_{j,t-\ell}$, which denotes the expected payoff of keeping the item. Finally, $\epsilon_{jt}^{A_\ell}$ is a seller-, time- and choice-specific idiosyncratic shock, which is assumed to be drawn independently across choices, sellers, and time. These shocks are privately observed by the seller and unknown to the econometrician.

Finally, the optimal auction for seller j at time t is simply given by $\pi_{jt}^{A} = \max_{\ell \in \mathbb{L}_t} \pi_{jt}^{A_\ell}$. Notice that the seller's auction choice set is time-dependent since some auction lengths become unavailable when the deadline is sufficiently close.

3.2.2 Posted prices

Posted prices are simpler mechanisms than auctions and, because sellers are assumed to make their decisions daily, they do not require a commitment restriction such as Assumption 7. However, since posted prices have no deadlines a seller has two value functions associated with them: one for creating a new posted price listing and another for an existing posted price. The value function from creating a new posted price listing on day t for seller j, which is analogous to expression (6), is given by

$$\pi_{jt}^{P} = \max_p \mathbb{E}_{N_{jt}^{P}} \left[p \Pr\left(V_{jt}^{(n:n)} \geq p\right) + \Pr\left(V_{jt}^{(n:n)} < p\right) \Pi_{j,t-1}(p) \middle| N_{jt}^{P} = n \right] - \kappa_j + \epsilon_{jt}^{P}$$
$$\equiv \tilde{\pi}_{jt}^{P} + \epsilon_{jt}^{P}. \tag{7}$$

The first term simply indicates that a sale is made if at least one buyer is willing to accept the posted price, which is equivalent to the probability that the highest valuation among all buyers who were matched to the listing exceeds the posted price. The dependence of the continuation value on the price chosen by the seller arises precisely because posted prices have no ending date. Thus, if a seller reaches day t with an existing posted price listing with price p', this seller has the option of doing nothing, which yields the following value function for an existing posted price:

$$\begin{aligned}\pi_{jt}^{P_e}(p') &= \mathbb{E}_{N_{jt}^P}\left[p' \Pr\left(V_{jt}^{(n:n)} \geq p'\right) + \Pr\left(V_{jt}^{(n:n)} < p'\right)\Pi_{j,t-1}(p')\Big| N_{jt}^P = n\right] + \epsilon_{jt}^{P_e} \\ &\equiv \tilde{\pi}_{jt}^{P_e}(p') + \epsilon_{jt}^{P_e}.\end{aligned} \tag{8}$$

Notice that a seller does not need to pay the listing cost κ_j in case no changes are made to the listing. This cost is meant to reflect the disutility a seller suffers from having to return to the platform and relist the tickets as well as choosing its features.

3.2.3 Sellers' outside options and optimal mechanism choice

Finally, seller j has an outside option given by

$$\pi_{jt}^O = \tilde{\pi}_{jt}^O + \epsilon_{jt}^O. \tag{9}$$

If chosen, the seller leaves the platform and does not return. It captures a series of possibilities that I do not observe in the data, such as the seller selling the tickets somewhere else, giving them away, or using them to go to the game herself.

Putting all value functions together and suppressing the prices in value functions of posted prices for ease of notation, seller j's choice on day t is simply given by

$$k_{jt}^* = \begin{cases} \arg\max_{k\in\{A,P\}}\{\pi_{jt}^A, \pi_{jt}^P\}, & \text{when } j \text{ arrives to the market} \\ \arg\max_{k\in\{A,P,O\}}\{\pi_{jt}^A, \pi_{jt}^P, \pi_{jt}^O\}, & \text{if } j \text{ does not have a posted price listing} \\ \arg\max_{k\in\{A,P,P_e,O\}}\{\pi_{jt}^A, \pi_{jt}^P, \pi_{jt}^{P_e}, \pi_{jt}^O\}, & \text{otherwise} \end{cases}. \tag{10}$$

Seller entry is taken as exogenous in this model, so at the moment of entry the outside option is not available to the seller. When the seller has an existing posted price listing at price p on day t, her choice will be given by $\max\{\pi_{jt}^A, \pi_{jt}^P, \pi_{jt}^{P_e}, \pi_{jt}^O\}$; if not, it will be given by $\max\{\pi^A_{jt}, \pi^P_{jt}, \pi^O_{jt}\}$. Hence, continuation values if seller j has to return at time t to make a new decision after a failed sale attempt is

$$\Pi_{jt} = \begin{cases} \mathbb{E}_{\epsilon_{jt}}\left[\max\{\pi^A_{jt}, \pi^P_{jt}, \pi^O_{jt}\}\right], & \text{if } j \text{ does not have a posted price listing} \\ \mathbb{E}_{\epsilon_{jt}}\left[\max\{\pi^A_{jt}, \pi^P_{jt}, \pi^{P_e}_{jt}, \pi^O_{jt}\}\right], & \text{otherwise} \end{cases} . \tag{11}$$

3.3 Equilibrium

To close the model I need to impose that buyers' and sellers' choices are consistent with one another, that is, I need to specify an equilibrium, whose definition is given below.

Definition 1. *An **equilibrium** is a set of distributions* $\left\{F^{A_t}_{R,T\Lambda,X}(\cdot), F^{P_t}_{P,\Lambda,X}(\cdot)\right\}_t$ *such that:*

1. *At each t, given $\left(F^{A_t}_{R,T\Lambda,X}(\cdot), F^{P_t}_{P,\Lambda,X}(\cdot)\right)$, arriving consumers choose which market to join according to equations (1), (2), (3), (4), and (5).*
2. *Given the evolution of market tightness, sellers solve the dynamic problem given by equations (6), (7), (8), (9), (10), and (11).*
3. *Market tightness, α^k_{jt}, is determined by which market buyers choose to enter and by which mechanism sellers choose to employ.*
4. *The remaining features of the distributions $\left(F^{A_t}_{R,T\Lambda,X}(\cdot), F^{P_t}_{P,\Lambda,X}(\cdot)\right)$ are determined by sellers' choices.*

This equilibrium notion relies on individual buyers and sellers being small relative to the entire market: the equilibrium distributions are taken as given by individual buyers and sellers, who further assume that these distributions are unaffected by their individual choices. However, when taken collectively their choices determine equilibrium distributions and outcomes. Finally, it is important to mention that establishing existence or uniqueness of such equilibrium in beyond the scope of this paper; instead, it is presented to close the model, and the estimation procedure does not require solving for it.

4 Estimation

I now describe the estimation procedure to recover the model's primitives. The procedure has three steps: first, the distributions of valuations are recovered from bidding data. Second, given these distributions, the parameters of buyer arrival processes are estimated separately for each selling mechanism. Third, given the aforementioned estimates, the distribution from which listing cost and outside option parameters are drawn are estimated by solving the sellers' problem by backward induction. All steps involve parametric assumptions, which are explicitly stated. I conclude by providing a brief discussion of identification.

4.1 Distributions of valuations

Distributions of valuations are recovered from bidding data. A well-known difficulty with the empirical analysis of online auctions is that the number of potential bidders is not observed by the econometrician, which prevents the use of order statistic inversion techniques discussed in, for example, Athey and Haile (2002). Thus, I follow the approach pioneered by Song (2004) and make use of multiple bids to recover the underlying distribution of valuations.

Under a symmetric IPV framework, the conditional density of the highest bid given the second highest is given by

$$g(v_1|v_2) = \frac{f(v_1)}{1 - F(v_2)}, \tag{12}$$

where $f(\cdot)$ and $F(\cdot)$ denote the pdf and cdf of the parent distribution, respectively. Importantly, this expression does not depend on the number of bidders. Even though the distribution $F(\cdot)$ is nonparametrically identified from this relation, I adopt a parametric approach. This is because of the limited number of observations on each day compounded with the curse of dimensionality with respect to the dimension of X_{jt}, and the fact that the estimated distributions are a key input to recover the remaining primitives of the model. I assume that valuations follow a Rayleigh distribution, with parameter $\sigma_{jt}^2 = \exp(X_{jt}'\mu)$.[6] I choose this distribution solely due to its simplicity, which facilitates estimation and yields

[6] A random variable, V, that follows a Rayleigh distribution with parameter σ^2 has the following probability density function: $f(v) = \frac{v}{\sigma^2} \exp\left\{-\frac{v^2}{2\sigma^2}\right\}$ and its cumulative distribution function is $F(v) =$

a closed-form expression for the optimal reserve price. I include in X_{jt} weekly intercepts and home team, away team, and season round indicators. I also categorize tickets into ten levels of quality and account for them using dummies.[7] This choice was motivated by the empirical patterns documented previously, which showed that the market drastically changes each week before the games. The vector of parameters μ is estimated via partial maximum likelihood (PMLE) based on expression (12).

4.2 Arrival processes

I approximate the arrival processes parameters, λ^k_{jt}, using the number of different users who clicked on listing j on day t. In other words, I treat this number as the realization of the draw from the Poisson distribution. While this is imperfect, it has the advantage of being much simpler than using the probability of a sale along with the estimates of the distributions of valuations to back out the implied arrival parameters.

More specifically, let C^k_{jt} denote the number of different users who click on listing j on day t, $Z_{jt} = (Z_j, t)$, and $\lambda^k_{jt} = \exp\{Z'_{jt}\lambda^k\}$. The vector λ^k can be estimated from the following Poisson regression model:

$$\begin{aligned} C^k_{jt} &= \Lambda^k_{jt} + \nu^k_{jt} \\ &= \lambda^k_{jt}\alpha^k_{jt} + \nu^k_{jt} \\ &= \exp\{Z'_{jt}\lambda^k\}\alpha^k_{jt} + \nu^k_{jt}, \end{aligned} \tag{13}$$

where α^k_{jt} is market k's tightness on day t, which is observed in the data, the vector Z_{jt} includes week indicators, and dummies for the quartiles of seller scores and ratings as well as types, and the error term ν^k_{jt} is such that $\mathbb{E}[\nu^k_{jt}|Z_{jt}, \alpha^k_{jt}] = 0$. The parameters λ^k are estimated via nonlinear least squares (NLLS) separately for auctions and posted prices.

4.3 Listing cost and outside option parameters

Having estimated the distributions of valuations and the parameters of the arrival processes, the remaining parameters to be estimated are the listing cost and outside option

$1 - \exp\left\{-\frac{v^2}{2\sigma^2}\right\}$.

[7] These levels are interactions between upper, club, and lower levels with sideline, corner, or end zone seats, as well as a VIP category.

parameters. I specify:

$$\tilde{\pi}^{O}_{jt} = \text{fv}_j \left(\psi_{1j}\mathbb{1}\{t < 8\} + \psi_{2j}\mathbb{1}\{7 < t < 15\} + \psi_{3j}\mathbb{1}\{14 < t < 22\} + \psi_{4j}\mathbb{1}\{t > 21\} \right), \quad (14)$$

where fv_j is the face value of j's tickets. The factor in the outside option varies across weeks to account for the fact that proximity to the deadline can affect sellers aggregately.

Given this parametrization, the sellers' problem is solved by backward induction, which is straightforward because they face a finite-horizon problem. Posted prices are chosen from a discretized grid, while the optimal reserve price, r, satisfies the usual equation given, for instance, by Myerson (1981). Finally, the term $\tilde{\pi}^{A_\ell}_{jt}$ is computed using numerical integration techniques, with details presented in Appendix C.3.

The estimation procedure works as follows. I assume that the elements of the vector of shocks ϵ_{jt} are drawn from a Type-1 Extreme Value (T1EV) distribution with location parameter $-\Gamma$ and scale parameter equal to 1, where Γ is the Euler-Mascheroni constant, independently distributed across time, alternatives, and sellers, which implies that all unobserved and persistent seller heterogeneity is captured by the listing cost and outside options. For any value of $\theta_j = (\kappa_j, \psi_{1j}, \psi_{2j}, \psi_{3j}, \psi_{4j})$, the sequence of mechanism and price choices made by a seller is unique. Let T_j be the set which contains all the days in which seller j is observed making active choices.[8] Given θ_j, j's individual likelihood is given by

$$l_j(\theta_j) = \prod_{t \in T_j} \left[\prod_{k \in \mathbb{L}_t} \left(\frac{\exp(\tilde{\pi}^k_{jt})}{\sum_{k' \in \mathbb{L}_t} \exp(\tilde{\pi}^{k'}_{jt})} \right)^{\mathbb{1}\{k_{jt}=k\}} \right]. \quad (15)$$

However, recall that the vector θ_j is unknown and seller-specific. Thus, I assume that $\theta_j \overset{iid}{\sim} H(\theta)$, so that the individual likelihood becomes

$$l_j(\theta) = \int_{\Theta} \prod_{t \in T_j} \left[\prod_{k \in \mathbb{L}_t} \left(\frac{\exp(\tilde{\pi}^k_{jt})}{\sum_{k' \in \mathbb{L}_t} \exp(\tilde{\pi}^{k'}_{jt})} \right)^{\mathbb{1}\{k_{jt}=k\}} \right] dH(\theta_j|\theta), \quad (16)$$

and the log-likelihood function of the data is given by $\mathcal{L}(\theta) = \sum_{j=1}^{J} \log[l_j(\theta)]$.[9]

[8] For example, if seller j enters ten days before the game, chooses a three-day auction that is not successful, followed by a posted-price listing in the next two days, and then exits the market without selling the tickets, then $T_j = \{5, 6, 7, 10\}$.

[9] It is important to note that this approach is not as efficient as possible because predicted prices are not matched to the observed ones.

To compute and maximize this likelihood the distribution $H(\theta)$ needs to be specified. I assume that

$$\log \begin{pmatrix} \kappa_j \\ \psi_{1j} \\ \psi_{2j} \\ \psi_{3j} \\ \psi_{4j} \end{pmatrix} \overset{iid}{\sim} N\left(\bar{\theta}, \Omega\right), \tag{17}$$

where Ω is a diagonal matrix. Rather than numerically solving the integral in (16), I employ the importance sampling simulation procedure proposed by Ackerberg (2009). For each observation j, I draw Q simulation draws from an importance sampling density $\iota(\cdot)$. Denoting each of these draws by θ_{jq} and letting $\varphi(\cdot)$ denote the probability density function of the lognormal distribution, the individual simulated likelihood is

$$\hat{l}_j(\bar{\theta}, \Omega) = \frac{1}{Q} \sum_{q=1}^{Q} \left\{ \prod_{t \in T_j} \left[\prod_{k \in \mathbb{L}_t} \left(\frac{\exp(\tilde{\pi}^k_{jt,q})}{\sum_{k' \in \mathbb{L}_t} \exp(\tilde{\pi}^{k'}_{jt,q})} \right)^{\mathbb{1}\{k_{jt,q}=k\}} \right] \right\} \frac{\varphi(\theta_{jq}|\bar{\theta}, \Omega)}{\iota(\theta_{jq})}, \tag{18}$$

where the subscript q also indicates that the model is solved and conditional choice probabilities are computed for each simulation draw. To estimate the parameters $\left(\bar{\theta}, \Omega\right)$ I maximize the simulated log-likelihood function $\hat{\mathcal{L}}(\bar{\theta}, \Omega) = \sum_{j=1}^{J} \log[\hat{l}_j(\bar{\theta}, \Omega)]$. I choose $\iota(\cdot)$ to be independent uniform distributions between zero and 1.5 for the outside option parameters and between zero and 25 for the listing costs. To guarantee that the resulting estimator is consistent and asymptotically normal it is necessary that the number of simulation draws increases faster than the squared root of the sample size,[10] so I set $Q = \lfloor J^{0.6} \rfloor = 372$, and I allow the set of parameters to vary across the three aforementioned seller types.

4.4 Identification

I now provide a brief discussion of identification. The distribution of buyer valuations is identified from bidding data alone because of the relationship established in equation (12). For this relationship to be true, bids must equal buyers' valuations, which requires several conditions to hold, more crucially Assumptions 1 and 3. Under these conditions, short-lived buyers are exogenously matched to listings, and those who are matched with auctions find it optimal to submit their valuations as bids provided that they are not lower

[10] See, for example, chapter 10 in Train (2009).

than the highest losing bid. The random matching and bid submission further imply that auctions with at least two submitted bids can be used to estimate the distribution of buyer valuations.

Perhaps the most critical condition for the aforementioned identification result to hold in the context of mechanism choice is Assumption 1(*ii*): buyers draw their valuations after they are matched to a listing. Crucially, this rules out buyers self-selecting into specific mechanism markets based on their valuations. This restriction is maintained to keep the model tractable despite its unlikelihood. An alternative approach would be to frame the buyer's problem as a multinomial discrete choice one, in which the buyer would choose between different listings by maximizing the expected payoff he would obtain from each alternative, and this choice set could be constructed based on clicking data. This would rely on a simple search model where buyers pay a fixed search cost and randomly draw a number of listings from the existing supply, which would correspond to the listings they clicked on. However, the difficulty in this approach would be handling the expected payoff from auctions since it would rely on the expected highest competing bid, which would be an equilibrium object.[11] One possibility would be to adopt a two-step approach, recovering the distribution of the highest bid directly from data in the first step and solving the buyer's multinomial discrete choice problem in the second step, but given the limited amount of bidding data the resulting estimates in all likelihood would be too inflexible and noisy. Given that this is only one of the features of the model, whose focus is actually on the supply side, I therefore chose a simpler modeling framework.

The parameters associated with the matching technology that brings buyers to listings are recovered based on equation (13), leveraging data on clicks, seller characteristics and observed market tightness. For this approach to be valid, several conditions need to hold. First, buyers and the matching technology need to operate as described in Section 3.1. Second, clicks need to accurately reflect matches, as addressed in the discussion of Assumption 1. Finally, the specification relies on an asymptotic approximation where the number of buyers and sellers is large. This condition is also present in the aforementioned equilibrium concept, in which buyers and sellers are sufficiently small in the market so that their actions do not directly alter equilibrium distributions. While this is satisfied for a fraction of the games in the sample, it might be a more stringent requirement for others.

[11] Determining the existence of such an equilibrium, even when seller choices are exogenous, is a complicated task. See Maslov (2020b) for an example.

Conditional on the the recovered distribution of buyer valuations and arrival processes parameters, the distribution of seller outside options and listing costs is identified by solving the seller's dynamic mechanism choice problem under the assumption that sellers behave as described in 3.2. The distribution of buyer valuations along with the Poisson arrival rates enables the computation of the instant expected payoff from running an auction with any given length and reserve price and from listing tickets as posted prices for any posted price level. Under the model's specification of forward-looking behavior, choices between auctions and posted prices, between auction lengths and between staying in the market instead of leaving recover the distribution of outside options and listing costs. It is important to note that this distribution is jointly identified by agents behaving as described in the model, by all parametric assumptions made and by the highlighted patterns in the data.

4.5 Limitations

Even though this estimation procedure recovers all structural parameters of interest in the context of the model, it does not leverage all the information available. As a consequence, the model should not be expected to accurately approximate the distributions of many of the features of the data.

While bids are used to recover the distributions of valuations, posted and reserve prices are obtained from the seller's optimization problem and are not matched to the observed prices in the data. As a consequence, to the extent that the modeling assumptions depart from what happens in reality these quantities should differ as well. However, given the dynamic structure of the model, it captures the patterns of falling prices as the deadlines approach. In addition, the model imposes restrictions on the matching procedure that are not reconciled with the observed market tightness in the data. The implied number of buyers and sellers implied by the estimates are unlikely to fit what is observed in the data, although the individual number of clicks each listing receives should.

Finally, the main patterns the model attempts to explain, namely seller exit and especially mechanism choices, are used to fit the data. Hence, these patterns should be reasonably approximated by the model.

5 Results

I now present the estimation results from the aforementioned procedure. First, results regarding the distributions of valuations are presented. Then, I display results for the arrival processes parameters. Finally, I show results for the distribution of seller-specific parameters.

5.1 Distributions of valuations

Following the procedure described above, Table 9 displays the estimates for the key parameters of the distributions of buyer valuations.

Table 9: Main parameters of distributions of valuations

Indicator	Estimate	*t*-statistic
Week of the game	6.83	29.63
Two weeks before game	7.13	30.65
Three weeks before game	7.14	30.48
Four weeks before game	7.16	27.3
Number of auctions	6,292	

Notes: Home team, away team, season round, and ticket type dummies are omitted for ease of exposition. Robust standard errors were used to construct the *t*-statistics.

Since the estimates are based on the Rayleigh distribution, interpretation is not direct. Instead, these estimates indicate the statistical significance of the results. Estimates of the dummies for home team, away team, season round, and ticket type are omitted for ease of exposition. All time indicators are highly significant, as the large values of the *t*-statistics indicate. The point estimates decrease as the deadline approaches, going from 7.16 four weeks before the game to 6.83 on the week of the game. This implies that the distributions of valuations display a first-order stochastic ordering over time. To illustrate this ordering, Figure 6 displays the distributions for valuations for VIP tickets of a hypothetical game between the Dallas Cowboys and the San Francisco 49ers in Dallas on the sixth week of the season.

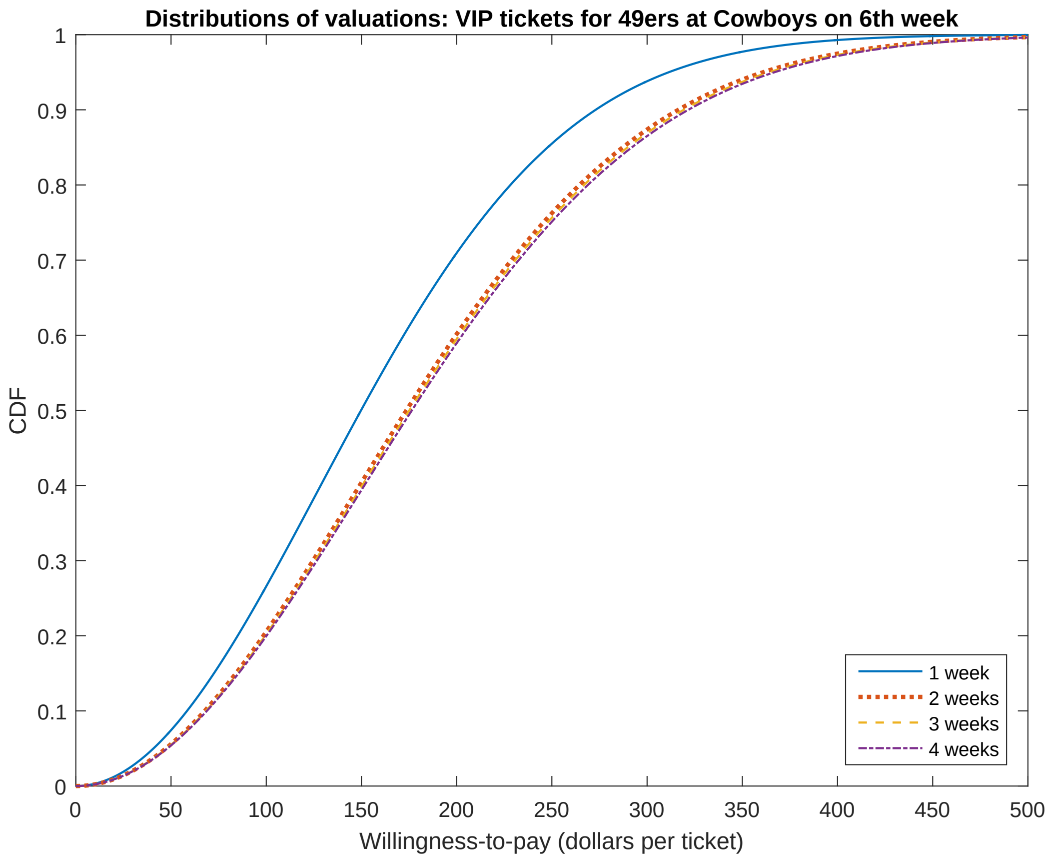

Notes: Figure shows Rayleigh distributions using the PMLE estimates of equation (12).

Figure 6

5.2 Arrival processes parameters

I now present estimates of the arrival processes parameters based on NLLS applied to equation (13). Parameter estimates are displayed in Table 10.

Due to the nonlinear specification of the Poisson parameters these coefficients cannot

Table 10: Estimates of Poisson arrival processes parameters

Parameter	Auctions	Posted prices
Intercept (week 1)	-0.47 (-29.61)	0.02 (0.88)
Intercept (week 2)	-0.55 (-34.52)	0.08 (3.43)
Intercept (week 3)	-0.54 (-31.06)	0.05 (1.76)
Intercept (week 4)	-0.55 (-27.02)	0.17 (4.28)
Medium seller	-0.08 (-5.41)	-0.21 (-8.17)
Big seller	-0.02 (-1.24)	-0.44 (-11.58)
Seller's score (2nd quartile)	-0.04 (-0.67)	0.14 (1.25)
Seller's score (3rd quartile)	-0.03 (-0.55)	0.07 (0.66)
Seller's score (4th quartile)	-0.08 (-1.44)	-0.06 (-0.51)
Seller's rating (2nd quartile)	0.46 (7.62)	0.15 (1.31)
Seller's rating (3rd quartile)	0.43 (7.12)	0.09 (0.78)
Seller's rating (4th quartile)	0.44 (7.63)	0.09 (0.84)

Notes: Table shows estimates of equation (13), ran separately for auctions and posted prices, with *t*-statistics displayed between parentheses. Robust standard errors were used to construct the *t*-statistics.

be interpreted directly, but the results indicate the statistical significance of the chosen variables as well as the overall correlations. Larger sellers attract less buyers for both auctions and posted prices, although the coefficient for large sellers using auctions is not significant, and sellers' scores are never significant. Sellers' ratings are only significant for auctions and better rated sellers attract more buyers, as expected. Finally, the overall baseline over time seems to indicate that posted prices attract more buyers than auctions as the coefficients associated with the former are positive and those with the latter are negative.

5.3 Distributions of listing costs and outside options

Based on the previous estimates, I estimate the distribution of seller-specific parameters following the method discussed in Subsection 4.3. Results are displayed in Table 11.

All coefficients are highly significant, which is consistent with the large heterogeneity in seller behavior displayed in the descriptive analysis. To better illustrate the results, Figure 7 plots all the distributions for seller parameters. Medium sellers have the lowest listing costs, followed by big and small sellers. In turn, large sellers have the lowest outside options, followed by medium and then small sellers. Outside options tend to increase as the deadline approaches.

These results are consistent with the overall profiles of seller types. In the sample large sellers are usually sports or tickets stores, which probably offer tickets for NFL games as a part of their business, but not as their central role. Hence, they do not intend to closely monitor the listings, which is reflected in their high listing costs. Furthermore, unlike individual sellers these stores do not have friends, co-workers, or neighbors to whom they can offer their tickets, and naturally cannot go to the game themselves. These factors are arguably why the larger the seller, the lower the outside option. The difference between small and medium sellers in regards to their listing costs comes from the fact that medium sellers offer several sets of tickets throughout the sample, which indicates they might be more likely to sell as an actual business activity instead of just offering tickets for games they cannot attend or are not interested in. Finally, outside options are higher on game week possibly because many sellers have the option of going to the games themselves.

Table 11: Estimates of sellers' listing cost and outside option parameters

	Seller type		
Mean	Small	Medium	Big
Listing cost (κ)	1.66 (67.09)	1.36 (68.05)	1.6 (29.16)
Outside option: 1 week (ψ_1)	-0.39 (-15.94)	-0.44 (-18.86)	-0.67 (-16.01)
Outside option: 2 weeks (ψ_2)	-0.78 (-16.97)	-0.93 (-26.6)	-1.02 (-13.6)
Outside option: 3 weeks (ψ_3)	-0.73 (-13.19)	-0.84 (-20.21)	-1.05 (-14.71)
Outside option: 4 weeks (ψ_4)	-0.73 (-8.55)	-0.76 (-15.24)	-0.79 (-11.65)
Variance	Small	Medium	Big
Listing cost (ω_κ^2)	1 (32.64)	1.27 (40.26)	1.39 (18.7)
Outside option: 1 week (ω_1^2)	0.49 (19.55)	0.48 (18.17)	0.77 (10.2)
Outside option: 2 weeks (ω_2^2)	0.41 (11.1)	0.44 (15.44)	0.58 (8.18)
Outside option: 3 weeks (ω_3^2)	0.35 (8.64)	0.38 (12.98)	0.55 (8.07)
Outside option: 4 weeks (ω_4^2)	0.31 (5.36)	0.31 (10.05)	0.36 (8.46)

Notes: Table shows estimates of the simulated log-likelihood function based on equation (18). Standard errors were obtained via nonparametric bootstrap based on 50 replications, and t-statistics constructed using them are displayed in parentheses.

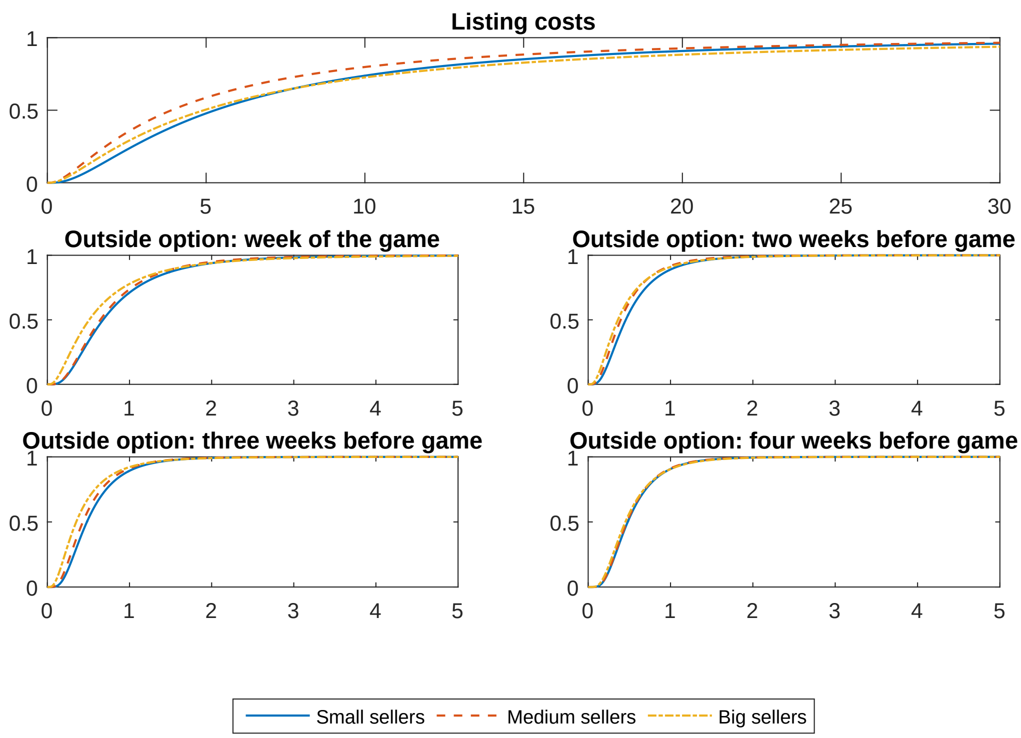

Notes: Figure shows lognormal distributions using the importance sampling simulated MLE estimates of equation (18).

Figure 7

6 Counterfactuals

Based on the estimates of the model's primitives I now conduct counterfactual exercises in which the menu of available mechanisms is altered. In particular, I investigate how removing all auctions or posted prices from the sellers' choice set impacts their decisions as well as transaction outcomes. To conduct these counterfactual exercises, I simulate 100 realizations of each game in the data under three types of scenarios: when all mechanisms are available, when only posted prices are, and when only auctions are. For each simulation I draw new shocks, listing costs, outside options, number of arriving buyers, and valuations of said buyers. I only consider situations in which only the auction lengths offered by eBay are available. In addition, given the assumption that sellers have perfect foresight over market tightness in both markets, for simulations in which sellers have access to all mechanisms I do not compute the aforementioned equilibrium, but rather use the sequence I observe in the data (as in Figure 5a). For the purposes of the counterfactual scenarios, I make the following assumption.

Assumption 8.
(i) The same sellers, with the same set of tickets, would have entered the market at the same point in time regardless of which mechanisms were available to them.
(ii) The overall market tightness would have followed the same trajectory in all three scenarios.

Assumption 8 deserves a few comments. Part 8(*i*) implies that seller entry was indeed random. However, it could be argued that since eBay was the main secondary market for tickets at the time that offered auctions as a listing mechanism, many would have chosen not to participate in it if auctions had not been available. In turn, part 8(*ii*) implies that since seller exit is endogenous, when the number of listings on the platform was higher (lower) than what is observed in the data, the number of buyers in the platform would increase (decrease) to keep market tightness constant. Notice that it also implies that I do not need to compute an equilibrium, since in a platform with just one type of mechanism buyers do not need to decide which market to join.

The first counterfactual I perform is to remove all auctions from the platform. This exercise was motivated by the fact that the use of online auctions has substantially decreased. As addressed by Einav et al. (2016) and Cullen and Farronato (2020), TaskRabbit began as an auction-only platform but since then has abandoned auctions altogether. The same phenomenon was studied by Wei and Lin (2017) and Huang (2020) in the context of Prosper.com. Within eBay, Einav et al. (2018) showed that sellers are moving away from

Table 12: Counterfactual results

Scenario	Object	Percent change			
		Mean	Median	10th %	90th %
Only posted prices	Pr(sale)	-84.27	-85.34	-89.93	-79.97
	$\mathbb{E}$ [payoff|sale]	-14.31	-12.46	-23.86	5.33
	$\mathbb{E}$ [payoff]	-87.37	-87.41	-91.62	-83.47
Only 1-,3-,5-,7-, and 10-day auctions	Pr(sale)	22.09	13.96	2.16	48.28
	$\mathbb{E}$ [payoff|sale]	-20.68	-19.22	-32.32	-11.4
	$\mathbb{E}$ [payoff]	-4.34	-8.02	-16.3	6.89
Only 1- and 3-day auctions	Pr(sale)	20.43	12.97	2.22	42.61
	$\mathbb{E}$ [payoff|sale]	-26.94	-26	-38.37	-16.76
	$\mathbb{E}$ [payoff]	-13.15	-15.13	-23.19	-6.15
Only 1- and 5-day auctions	Pr(sale)	21.27	13.6	2.18	45.18
	$\mathbb{E}$ [payoff|sale]	-25.35	-24.6	-36.92	-15.42
	$\mathbb{E}$ [payoff]	-10.62	-13.36	-21.26	-1.98
Only 1- and 7-day auctions	Pr(sale)	19.41	11.62	2.07	42.35
	$\mathbb{E}$ [payoff|sale]	-28.89	-27.92	-40.1	-19.4
	$\mathbb{E}$ [payoff]	-16.04	-18.23	-26.81	-8.06

Notes: Counterfactual results are averages across 100 simulations and changes are compared to the current options on eBay. Average sale prices are calculated with respect to the face values of the tickets.

auctions towards other mechanisms. These changes bring into question whether auctions can be helpful to sellers, as theory predicts. Results are given in the first row of Table 12. Eliminating auctions from the platform has a strong negative effect on sellers: expected transaction prices fall by a little more than 14% while the probability of sale falls by more than 84%, yielding a total decrease of more than 87% on sellers' expected payoffs. Without auctions the value sellers accrue from being on the platform decreases substantially, leading them to be more likely to take their outside option and leave after a failed attempt to sale. Posted prices are only effective when the deadline is sufficiently close because of buyers' relative reluctance to participate in auctions.

The natural comparison to a market without auctions is an all-auction platform. To do this I eliminate the possibility of using posted prices and keep the original auction options, that is, the possibility of choosing between one, three, five, seven, and ten days for auction length. Results for this counterfactual exercise are displayed in the second row of Table 12. In a world without posted prices sellers become more likely to sell as the probability of sales increases by 22.09%. However, transaction prices fall by more than 20%, yielding a total negative effect on overall expected revenues of more than 4%. This effect is driven by the fact that eliminating posted prices also reduces the sellers' value from being on the platform, but not enough for them to leave. However, this reduction does lead them to pick lower reserve prices, which yields more sales at lower prices. Interestingly, buyers would unconditionally benefit from such a scenario as not only would they be more likely to purchase tickets, but they would also pay lower prices. It is also interesting to note that the differences between the first two rows of Table 12 are somewhat reminiscent of the contrasts displayed in Figures 5c and 5d: given that auctions are much more likely to convert than posted prices, it is not surprising that an all-auction platform would yield more sales than one with only posted prices. However, posted prices do yield higher transaction prices, which explains why the decrease in the sellers' continuation value has a higher impact on transaction prices on an all-auction platform than on one with only posted prices.

Intuitively having more options is beneficial, which could explain why reducing the set of mechanisms sellers have access to always diminishes their expected payoffs. However, this logic can be misleading as the forms of competition between and within selling mechanisms are different. Despite an overall negative effect on expected revenues, an auction-only platform does enhance the probability of a sale being made, and the fact that the decrease in the expected transaction price more than compensates this increased likelihood of selling is a consequence of the model estimates and not of the model itself. In other words, it could have been the case that the higher probability of sales was enough to offset the decrease in expected transaction prices, actually leading to an overall increase in expected revenues. Indeed, the 90th percentile of the percent change in expected payoffs from an all-auction platform is an actual increase of almost 7%, indicating that for same games an auction-only platform would be beneficial to sellers.

It could be the case that the discrepancy between the overall losses in expected revenues from an auction-only and a posted price-only platform is solely a consequence of the fact that the numbers of options eliminated from the sellers' choice set are different: eliminating posted prices amounts to taking out two options from the sellers' menu, creating a new posted price listing or keeping an existing one, while eliminating all auctions implies depriving sellers of five of their original options. Since I am assuming a T1EV distribution for the choice-specific error terms, a concern is that the magnitudes of decreases are therefore an artifact of this distributional assumption. To investigate whether this is the case I consider additional auction-only counterfactuals with fewer available auction lengths. In particular, I always maintain one-day auctions but only allow for one additional length, three, five, or seven days, with results displayed on rows three, four, and five of Table 12. I always keep the one-day auction as it is the most flexible auction alternative for sellers, and I always add just one additional alternative to make the comparison with the two alternatives afforded by posted prices as close as possible. I do not consider a scenario with only one- and ten-day auctions because the longest alternative is rarely chosen in the data and in the simulations.

Even though having more options for auction length always benefits sellers, the qualitative results with a different set of available auction lengths are always remarkably similar: the probability of sales always increases (between 19.41% and 21.27%) but the decrease in expected transaction prices always more than compensates it (between 25.35% and 28.89%), leading to an overall decrease in expected revenues (between 10.62% and 16.04%). However, this overall decrease is always much smaller than the one resulting from eliminating all auctions instead of posted price and a subset of auctions, suggesting that this result is not due to the T1EV assumption. One possibility is that the estimates of arrival rates and seller parameters somehow capture the fact that auctions are precisely what attracts users to eBay since it is the main resale market for tickets that makes use of this selling mechanism.

7 Conclusion

This study examines the impact of the availability of different selling mechanisms on a perishable good market. It leverages data on NFL tickets offered on eBay to estimate a structural model in which heterogeneous, forward-looking sellers choose between the available mechanisms and their features, while short-lived buyers only decide which market to join, auctions or posted prices. The model is informed by a descriptive analysis of the data, which in itself is an additional contribution since parts of them are novel, such as the availability of clickstream data by potential buyers.

To assess the role of selling mechanisms, I utilize the estimates of the aforementioned model to conduct counterfactual exercises in which the menu of mechanisms available to sellers changes. I find that sellers' expected revenues would on average decrease by 87.37% if auctions were completely removed from their choice set. In turn, removing posted prices would only decrease average expected revenues by 4.34%. This demonstrates that while sellers benefit from the availability of a menu of mechanisms, most of its value comes from auctions. On the other hand, buyers would benefit from an auction-only platform as the expected number of transactions would increase but the expected transaction prices would decrease.

Put together, these findings have non-trivial consequences for platform design. In this setting, the platform as a two-sided market has to cater to both sides of the market to sustain market operation. These results indicate that there can exist a potential tension between the interests of the two sides of the market regarding which selling mechanisms the platform should provide. However, it is important to keep in mind that these results were obtained under a set of assumptions. A perishable good market with the coexistence of different selling mechanisms is a complex environment, which led me to make a series of important simplifications, such as assuming a less sophisticated demand side and disregarding explicit competition between sellers. Richer data could enable the specification and estimation of a more intricate model that could possibly yield different results. Therefore, revisiting this question under such circumstances could be an important exercise.

In addition, the estimated positive effect of the availability of auctions to both sides of the market is noteworthy as, in reality, online platforms are generally moving away from this selling mechanism. Nevertheless, these two facts are reconcilable as this study addresses only a particular type of perishable good, event tickets, while the platforms that moved away from auctions focused on goods and services of a different nature. Since improving markets for perishable goods is a continuing effort, a pertinent avenue for future research is establishing whether auctions have generally positive effects on this type of market.

Acknowledgements

I am hugely indebted to my advisors, Harikesh Nair, Brad Larsen, and Han Hong for constant guidance and support. Several conversations with Isa Chaves, Liran Einav, Pedro Gardete, Jon Levin, Guido Martirena, and Andy Skrzypacz substantially improved this study. I am especially thankful to Tom Blake and Dimitriy Masterov for great aid with the data and several helpful discussions throughout this project. I have also benefited from comments by Lanier Benkard, Vivek Bhattacharya, Tim Bresnahan, Matt Gentzkow, Brett Gordon, Wes Hartmann, Ken Hendricks, Kevin Hutchinson, Jim Lattin, Gordon Leslie, Sarah Moshary, Paulo Somaini, Nico Pierri, Navdeep Sahni, Stephan Seiler, Anna Tuchman, Milena Wittwer, and Frank Wolak, as well as audiences at several schools and conferences. Finally, I thank the editor, Günter Hitsch, and two anonymous referees for their detailed comments and feedback. Part of this work was conducted while I was a non-paid consultant to the economics team at eBay Research. Financial support from the Stanford Institute for Economic Policy Research through a dissertation fellowship is gratefully acknowledged. All mistakes are mine.

References

Ackerberg, D. A. (2009). A new use of importance sampling to reduce computational burden in simulation estimation. *Quantitative Marketing and Economics*, 7(4):343–376.

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

Backus, M. and Lewis, G. (2020). Dynamic demand estimation in auction markets. Working paper, Columbia University.

Bauner, C. (2015). Mechanism choice and the buy-it-now auction: A structural model of competing buyers and sellers. *International Journal of Industrial Organization*, 38:19–31.

Bhave, A. and Budish, E. (2018). Primary-market auctions for event tickets: Eliminating the rents of "Bob the Broker". Working paper, University of Chicago.

Board, S. and Skrzypacz, A. (2016). Revenue management with forward-looking buyers. *Journal of Political Economy*, 124(4):1046–1087.

Bodoh-Creed, A. L., Boehnke, J., and Hickman, B. (2021). How efficient are decentralized auction platforms? *Review of Economic Studies*, 88(1):91–125.

Caldentey, R. and Vulcano, G. (2007). Online auction and list price revenue management. *Management Science*, 53(5):795–813.

Coey, D., Larsen, B. J., and Platt, B. C. (2020). Discounts and deadlines in consumer search. *American Economic Review*, 110(12):3748–3785.

Courty, P. (2003). Some economics of ticket resale. *Journal of Economic Perspectives*, 17(2):85–97.

Cullen, Z. and Farronato, C. (2020). Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms. *Management Science*, forthcoming.

Eeckhout, J. and Kircher, P. (2010). Sorting versus screening: Search frictions and competing mechanisms. *Journal of Economic Theory*, 145(4):1354–1385.

Einav, L., Farronato, C., and Levin, J. (2016). Peer-to-peer markets. *Annual Review of Economics*, 8:615–635.

Einav, L., Farronato, C., Levin, J., and Sundaresan, N. (2018). Auctions versus posted prices in online markets. *Journal of Political Economy*, 126(1):178–215.

Einav, L., Kuchler, T., Levin, J., and Sundaresan, N. (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7(2):215–247.

Etzion, H. and Moore, S. (2013). Managing online sales with posted price and open-bid auctions. *Decision Support Systems*, 54(3):1327–1339.

Etzion, H., Pinker, E., and Seidmann, A. (2006). Analyzing the simultaneous use of auctions and posted prices for online selling. *Management Science*, 8(1):68–91.

Goettler, R. L. and Gordon, B. R. (2011). Does AMD spur Intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200.

Gowrisankaran, G. and Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, 120(6):1173–1219.

Hammond, R. G. (2010). Comparing revenues from auctions and posted prices. *International Journal of Industrial Organization*, 28(1):1–9.

Hammond, R. G. (2013). A structural model of competing sellers: Auctions and posted prices. *European Economic Review*, 60:52–68.

Hendricks, K. and Sorensen, A. (2018). Dynamics and efficiency in decentralized online auction markets. Working paper, University of Wisconsin.

Huang, G. (2020). Selling-mechanism design for peer-to-peer lending platforms: Multiunit uniform-price dynamic auctions vs. posted price. Working paper, Purdue University.

Hummel, P. (2015). Simultaneous use of auctions and posted prices. *European Economic Review*, 78:269–284.

Julien, B., Kennes, J., and King, I. P. (2001). Auctions and posted prices in directed search equilibrium. *Topics in Macroeconomics*, 1(1):1–16.

Kultti, K. (1999). Equivalence of auctions and posted prices. *Games and Economic Behavior*, 27(1):106–113.

Lee, C.-Y., Roberts, J. W., and Sweeting, A. (2012). Competition and dynamic pricing in a perishable goods market. Mimeo, Duke University.

Leslie, P. and Sorensen, A. (2014). Resale and rent-seeking: An application to ticket markets. *Review of Economic Studies*, 81(1):266–300.

Maslov, A. (2020a). Competition in online markets with auctions and posted prices. Working paper, Central Michigan University.

Maslov, A. (2020b). A note on buyers' behavior in auctions with an outside option. *Games*, 11(3):26–34.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Myerson, R. B. (1998). Population uncertainty and Poisson games. *International Journal of Game Theory*, 27(3):375–392.

Nair, H. S. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games. *Quantitative Marketing and Economics*, 5(3):239–292.

Rogerson, R., Shimer, R., and Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988.

Selcuk, C. (2017). Auctions vs. fixed pricing: Competing for budget constrained buyers. *Games and Economic Behavior*, 103:262–285.

Song, U. (2004). Nonparametric estimation of an eBay auction model with an unknown number of bidders. Mimeo, University of British Columbia.

Sun, D. (2008). Dual mechanism for an online retailer. *European Journal of Operational Research*, 187(3):903–921.

Sweeting, A. (2012). Dynamic pricing behavior in perishable good markets: Evidence from secondary markets for Major League Baseball tickets. *Journal of Political Economy*, 120(6):1133–1172.

Sweeting, A. (2013). Auctions vs. fixed prices vs. both: Mechanism choices for perishable goods. Mimeo, Duke University.

Sweeting, A. (2015). A model of non-stationary dynamic price competition with an application to market design. Mimeo, University of Maryland.

Sweeting, A. and Sweeney, K. (2015). Staggered vs. simultaneous price setting with an application to an online market. Mimeo, University of Maryland.

Talluri, K. and van Ryzin, G. (2004). *The Theory and Practice of Revenue Management*. Springer.

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press.

Wang, R. (1993). Auctions versus posted-price selling. *American Economic Review*, 83(4):838–851.

Wei, Z. and Lin, M. (2017). Market mechanisms in online peer-to-peer lending. *Management Science*, 63(12):4236–4257.

Zeithammer, R. and Adams, C. (2010). The sealed-bid abstraction in online auctions. *Marketing Science*, 29(6):964–987.

Zeithammer, R. and Liu, P. (2006). When is auctioning preferred to posting a fixed selling price? Mimeo, University of Chicago.

Zhang, H. (2020). Prices versus auctions in large markets. *Economic Theory*, forthcoming.

Ziegler, A. and Lazear, E. P. (2003). The dominance of retail stores. NBER working paper 9795.

Appendix

A Sample construction

I describe here in detail the procedure I used to create the chains of listings described in Section 2. The key information used to create the chains consists of the number of tickets being offered, the game they corresponded to, and the section and row where the seats were located. This information is key for two reasons. First, the linking process to track the same set of tickets over time is based exclusively on them.[12] Second, these variables allow me to identify the price of these tickets on the primary market, which yields a measure of their quality.

The first step is obtaining information on the game corresponding to a given listing, which is often available in a standardized fashion. When this is not the case, I obtained this information from the title or subtitle of the listings. When these are not informative, the dates in which listings were created by the sellers are ordered to potentially fill in this missing information. I also make use of this procedure to correct listings for tickets that were created after the games they corresponded to had taken place, which were usually instances in which the seller corrected the information shortly after. When this was not the case, sellers just removed the listing within a few days indicating that they were erroneous and possibly the result of automatic re-listing.

Of the valid listings with information regarding the game to which the tickets corresponded, around 97% also had information on number of tickets and section and row where the tickets were located. To fill in the missing information I use the listing's title or subtitle. I also verify whether sellers had offered multiple listings for the same team at the same location and whether the listings with missing information were created and terminated in between listings with complete information. This procedure was also useful to correct instances in which the information was erroneous, either because the section and row numbers were exchanged or because the information did not conform with what was reported on the title or subtitle.

With this information in hand I define as potential chains of listings combinations of different seller-game-section-row quadruples. I then identify instances in which list-

[12] Ideally the process would be based on the seat numbers being offered as they would make the creation of chains trivial. Unfortunately this information was rarely available, a difficulty also faced by Leslie and Sorensen (2014).

ings within the same chain are created before the previous one was over. These cases are inspected and classified into five scenarios. First, multiple chains at the same location offered by the same seller. This is done based on information on the titles, complete sales, and other chains by the same seller. Cases in which the seller creates two identical listings across all dimensions at virtually the same time are assumed to be for different tickets. The second scenario is reorganization of quantities, or rebundling. For example, turning a single listing for four tickets into two listings for two tickets each. Third, listings which were removed within a day and recreated shortly after are assumed to be mistakes and deleted. The final two scenarios concern listing the same set of tickets more than once concurrently, which I call doublelisting.

I classify doublelisting scenarios into two cases. The first is separation across quantities: for example, having a listing for four tickets and, at the same time, two separate listings for two tickets each. This is again cross-checked with the seller's history of listings and their outcomes. Returning to the example from the previous paragraph, if both a two-ticket and four-ticket listings are sold then they were for different sets of tickets, while if the four-ticket one is sold and the other two two-ticket listings are then removed from the website it suggests that the same tickets were listed twice. Finally, the second case consists of listing the same set of tickets through different mechanisms. I verified these cases according to the same procedure that I employed in the previous case.

At the end of this procedure I obtain a sample of 38,520 sets of tickets, which were offered across 78,863 listings. However, the analysis will be restricted to activity within four weeks of a game. This restriction is not extreme: in this period, almost 61% of tickets were introduced to the market, more than 70% of tickets were available at the website, and more than 67% of the transactions observed in the data took place. The resulting sample contains 27,047 sets of tickets across 43,215 listings. This subsample is used to construct the measures of market tightness and of potential demand on the platform at every point in time. To estimate the model, I drop the 293 listings, spread across 226 chains, that do not have information on the number of tickets, type of tickets, or section or row where tickets were located. I further restrict the sample to tickets that were always offered as a pair and were never rebundled or doublelisted. Therefore, the final sample is smaller, containing 19,174 pairs of tickets over 28,257 listings. Nevertheless, pairs are the most common bundle offered (more than 74% of listings).

B More mechanisms

I now document patterns relative to mechanisms that were not explicitly considered in the main text: auctions with an immediate purchase option, also known as buy-it-now (BIN) or hybrid auctions, and bargaining-enabled posted prices. First, Figure B.1 shows which among the four options sellers choose across time. Hybrid auctions are not as popular as regular auctions until the week of the game, possibly because the additional flexibility they yield becomes more attractive closer to the deadline, when buyers seem to be relatively less willing to participate in auctions. In turn, despite being the default option for posted prices, bargaining-enabled listings only become slightly more prevalent within two weeks of the game, possibly due to their additional flexibility akin to hybrid auctions.

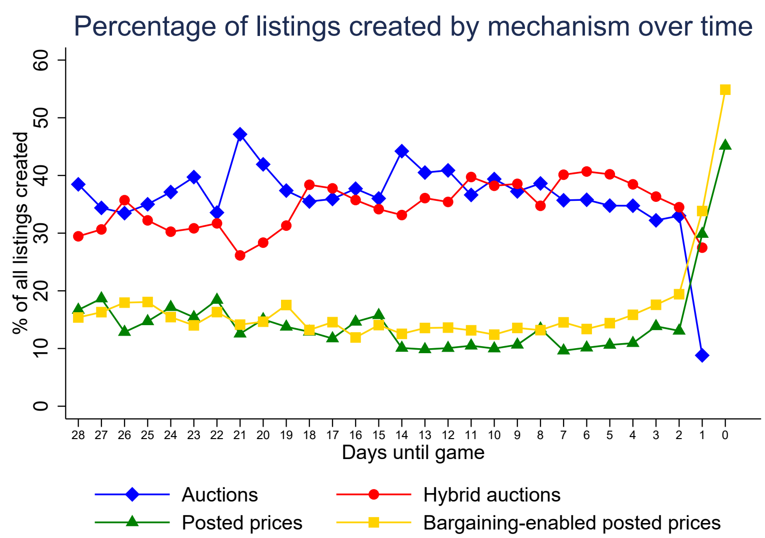

Notes: Quantities displayed on the vertical axis are averaged over all listings in the sample.

Figure B.1

A different consideration is if and how prices are chosen differently across these mechanisms. Figure B.2 displays the choice of start prices for regular and hybrid auctions, posted prices with and without the bargaining option, and buy prices for hybrid auctions. It is interesting to note that hybrid auctions consistently have higher start prices than regular ones, which could be a consequence of sellers with higher outside options self-selecting into the hybrid format. Moreover, posted prices with a bargaining option are consistently higher than those without it, possibly because sellers anticipate a potential negotiation that would likely reduce the final agreed upon price. Finally, it is interesting to note that buy prices in hybrid auctions are often higher than posted prices regardless of whether bargaining is available to buyers.

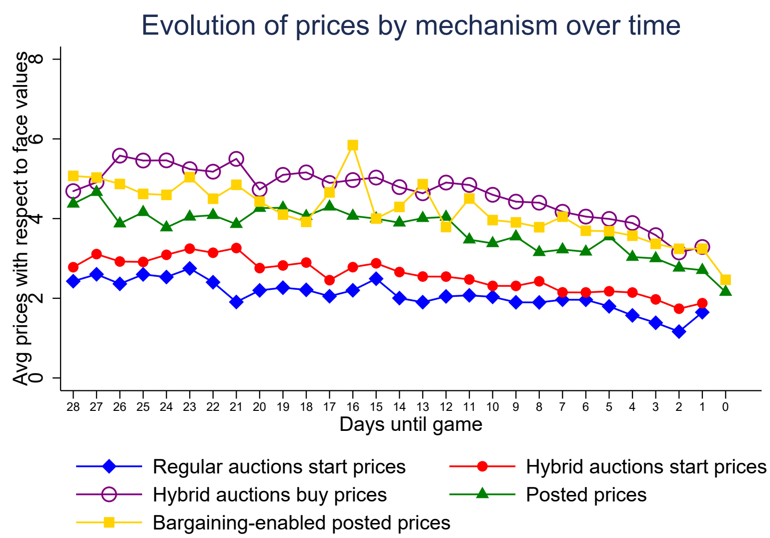

Notes: Quantities displayed on the vertical axis are averaged over all listings in the sample.

Figure B.2

A final consideration is whether buyers make use of this richer set of mechanisms. First, it is interesting to note that despite being slightly more prevalent than regular auctions, hybrid auctions are less likely to convert, possibly because of higher start and buy prices, as displayed in Figure B.2. In addition, just a little more than 15% of successful hybrid auctions were actually sold via buy prices, in part because when the reserve price is met the buy-it-now option goes away. Similarly, bargaining-enabled posted prices are more commonly used than regular posted prices, but have a lower conversion rate. Furthermore, even though almost half of the bargaining-enabled listings were involved in negotiations at some point, less than 30% of these were sold at a negotiated price. These numbers show that abstracting from these more detailed, hybrid mechanisms is not a dramatic simplification.

Table B.1: Distribution of listings across mechanisms II

Type	Quantity	Sold	Sold via buy price	
Auctions	10,050	4,580	–	
Hybrid auctions	10,087	3,802	588	
Type	Quantity	Sold	Bargained for	Sold via bargaining
Posted prices	3,661	1,020	–	–
Posted prices with bargaining	4,459	1,155	2,321	694

Notes: Table displays quantities of the final sample described in Section 2.

C Derivation of expected utility and profit functions

I now derive the expected utility and profit functions from Section 3.

C.1 Expected utility from auctions

For ease of notation I will ignore the subscripts. Assume that a buyer with original valuation v is matched on day t with an auction with reserve price r, end date τ, Poisson arrival rate λ, and characteristics x. Conditional on $v \geq r$ and on the number of opposing bidders being n, the buyer's expected utility from the auction is:

$$\begin{aligned}
U_{A_t}(v,r,\tau,n,x) &= \Pr(V^{(n:n)} < v|x,\tau)\left(v - \mathbb{E}\left[\max\left\{V^{(n:n)},r\right\}\middle|V^{(n:n)} < v,x,\tau\right]\right) \\
&= F_V(v|x,\tau)^n v - \int_0^v \max\{u,r\}\, nF_V(u|x,\tau)^{n-1} f_V(u|x,\tau)du \\
&= F_V(v|x,\tau)^n v - \int_0^r rnF_V(u|x,\tau)^{n-1} f_V(u|x,\tau)du \\
&\quad - \int_r^v unF_V(u|x,\tau)^{n-1} f_V(u|x,\tau)du \\
&= F_V(v|x,\tau)^n v - rF_V(r|x,\tau)^n - \left[uF_V(u|x,\tau)^n|_r^v - \int_r^v F_V(u|x,\tau)^n du\right] \\
&= \int_r^v F_V(u|x,\tau)^n du.
\end{aligned}$$

The number of opposing bidders, n, is unknown to the buyer. As demonstrated by Myerson (1998), the Poisson distribution yields environmental equivalence: the number of opposing bidders from a buyer's perspective will follow the same distribution as the one guiding the overall arrival of bidders. Hence, summing over all realizations of n:

$$\begin{aligned}
U_{A_t}(v,r,\tau,\lambda,x) &= \sum_{n=0}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \int_r^v F_V(u|x,\tau)^n du \\
&= \int_r^v \sum_{n=0}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} F_V(u|x,\tau)^n du \\
&= \int_r^v e^{-\lambda[1-F_V(u|x,\tau)]} du.
\end{aligned}$$

C.2 Expected utility from posted prices

Now assume that buyer i with valuation v is matched with a posted price with price p such that $v \geq p$ and arrival rate λ. If the number of opposing buyers is zero, a purchase is made. If there is one opposing buyer, with 50% probability i is called first and makes a purchase and with 50% probability i is called second and purchases because the other buyer's valuation is below p. Thus, applying this to all realizations of n yields:

$$\begin{aligned}
U_{P_t}(v,p,\lambda,x) &= (v-p)\sum_{n=0}^{\infty}\frac{\lambda^n e^{-\lambda}}{n!}\frac{1}{n+1}\sum_{l=0}^{n}F_V(p|x,t)^l \\
&= (v-p)\sum_{n=0}^{\infty}\frac{\lambda^n e^{-\lambda}}{(n+1)!}\left(\frac{1-F_V(p|x,t)^{n+1}}{1-F_V(p|x,t)}\right) \\
&= \frac{(v-p)}{1-F_V(p|x,t)}\left\{\frac{1}{\lambda}\sum_{n=1}^{\infty}\frac{\lambda^n e^{-\lambda}}{n!}-\frac{e^{-\lambda[1-F_V(p|x,t)]}}{\lambda}\right. \\
&\qquad\qquad \left.\times\sum_{n=1}^{\infty}\frac{[\lambda F_V(p|x,t)]^n\, e^{-\lambda F_V(p|x,t)}}{n!}\right\} \\
&= \frac{(v-p)}{\lambda\left[1-F_V(p|x,t)\right]}\left\{\left(1-e^{-\lambda}\right)-e^{-\lambda[1-F_V(p|x,t)]}\left(1-e^{-\lambda F_V(p|x,t)}\right)\right\} \\
&= \frac{(v-p)\left(1-e^{-\lambda[1-F_V(p|x,t)]}\right)}{\lambda\left[1-F_V(p|x,t)\right]}.
\end{aligned}$$

C.3 Computing choice-specific value function for auctions

I now describe how the choice-specific value functions for auctions ($\tilde{\pi}_{jt}^{A_\ell}$) were computed. To ease notation I will drop the subscripts for seller (j), date when the auction started (t), and auction length (ℓ), as well as the superscript indicating that an auction was chosen ($k = A$).

To calculate $\mathbb{E}\left[\max\{V^{(N-1:N)}, r\}\right]$, I first re-write the seller's payoff using analogous equivalences from static auction theory. First, for any number of bidders n, any reserve price r, and any seller continuation value, π_0, it follows that revenue can be expressed as $\max\{V^{(n-1:n)}, r\} + (\pi_0 - r)\,\mathbb{1}\{V^{(n:n)} < r\}$. Hence, the expected payoff from an auction is given by:

$$\pi^A = \mathbb{E}_N\left[\max\{V^{(n-1:n)}, r\}\middle|N=n\right] + (\pi_0 - r)\,Pr_N\left(V^{(n:n)} < r\middle|N=n\right). \qquad \text{(C.1)}$$

First, I compute the second term in the right-hand side of (C.1). It follows that, if bidders arrive according to a Poisson distribution with parameter λ,

$$\begin{aligned}(\pi_0 - r)\, Pr_N\left(V^{(n:n)} < r \middle| N = n\right) &= (\pi_0 - r) \sum_{n=0}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} F_V(r)^n \\ &= (\pi_0 - r)\, e^{-[1-F_V(r)]\lambda}. \end{aligned} \tag{C.2}$$

To compute the first term in the right-hand side of (C.1) first note that if $N \leq 1$, then $r \geq V^{N-1:N}$ with probability one. For these two values the expression simply is:

$$\sum_{n=0}^{1} \frac{\lambda^n e^{-\lambda}}{n!} r = e^{-\lambda} r\,(1+\lambda)\,. \tag{C.3}$$

Finally, consider now the case when $N > 1$. In particular, for any n and r it follows that:

$$\begin{aligned}\mathbb{E}\left[\max\left\{V^{n-1:n}, r\right\}\middle| N = n\right] &= \int_0^{\infty} \max\{u, r\} f_{n-1:n}(u) du \\ &= \int_0^{r} r f_{n-1:n}(u) du + \int_r^{\infty} u f_{n-1:n}(u) du. \end{aligned} \tag{C.4}$$

The two terms in (C.4) are now evaluated separately, beginning with the first in the right-hand side. Remember that due to properties of order statistics it follows that $F_{n-1:n}(r) = nF(r)^{n-1} - (n-1)F(r)^n$. Thus,

$$\begin{aligned}\mathbb{E}_N\left[rF_{n-1:n}(r)|N = n\right] &= r \sum_{n=2}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \left[nF_V(r)^{n-1} - (n-1)F_V(r)^n\right] \\ &= r\left\{\lambda\left[1 - F_V(r)\right] \sum_{n=2}^{\infty} \frac{\left[\lambda F_V(r)\right]^{n-1} e^{-\lambda}}{(n-1)!} + \sum_{n=2}^{\infty} \frac{\left[\lambda F_V(r)\right]^{n} e^{-\lambda}}{n!}\right\} \\ &= re^{-\lambda[1-F_V(r)]}\left\{\lambda\left[1 - F_V(r)\right]\left(1 - e^{-\lambda F_V(r)}\right)\right. \\ &\qquad \left. +1 - \left[1 + \lambda F_V(r)\right] e^{-\lambda F_V(r)}\right\} \\ &= re^{-\lambda[1-F_V(r)]}\left\{1 + \lambda\left[1 - F_V(r)\right]\right\} - re^{-\lambda}\,(1+\lambda)\,. \end{aligned} \tag{C.5}$$

Now I compute the expectation over N of the second term in the RHS of (C.4):

$$\begin{aligned}
\sum_{n=2}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \int_r^{\infty} y f_{n:1:n}(y) dy &= \int_r^{\infty} y \left[\sum_{n=2}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} f_{n-1:n}(y) \right] dy \\
&= \int_r^{\infty} y \left[\sum_{n=2}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \left[n f_{n-1:n-1}(y) - (n-1) f_{n:n}(y) \right] \right] dy \\
&= \int_r^{\infty} y \left[\sum_{n=2}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \left[n(n-1) F_V(y)^{n-2} f_V(y) - n(n-1) F_V(y)^{n-1} f_V(y) \right] \right] dy \\
&= \lambda^2 \int_r^{\infty} y e^{-\lambda[1-F_V(y)]} [1 - F_V(y)] f_V(y) \left[\sum_{n=2}^{\infty} \frac{[\lambda F_V(y)]^{n-2} e^{-\lambda[1-F_V(y)]}}{(n-2)!} \right] dy \\
&= \lambda^2 \int_r^{\infty} y e^{-\lambda[1-F_V(y)]} [1 - F_V(y)] f_V(y) dy \\
&= \int_0^{\lambda[1-F_V(r)]} x e^{-x} F_V^{-1} \left(1 - \frac{x}{\lambda} \right) dx \\
&= \int_0^{\lambda e^{-\frac{r^2}{2\sigma^2}}} x e^{-x} \sqrt{2\sigma^2 \log\left(\frac{\lambda}{x}\right)} dx \equiv \xi,
\end{aligned} \tag{C.6}$$

where the first equality follows from Fubini's theorem, the second from properties of order statistics, the penultimate from a change of variables in which $x = \lambda[1 - F_V(y)]$, and the last from the Rayleigh distribution which was assumed. The last integral, defined as ξ, is solved via Gauss-Chebyshev quadrature using ten nodes.

Plugging (C.5) and (C.6) in (C.4) yields:

$$\mathbb{E}_N \left[\max \left\{ V^{N-1:N}, r \right\} \right] = r e^{-\lambda[1-F_V(r)]} \left\{ 1 + \lambda \left[1 - F_V(r) \right] \right\} - r e^{-\lambda} \left(1 + \lambda \right) + \xi, \tag{C.7}$$

and plugging (C.2), (C.3), and (C.7) in (C.1) finally yields:

$$\pi^A = r\lambda[1 - F_V(r)] e^{-\lambda[1-F_V(r)]} + \xi + e^{-\lambda[1-F_V(r)]} \pi_0. \tag{C.8}$$