NET Institute*

www.NETinst.org

Working Paper #15-03

September 2015

**A Model of Non-Stationary Dynamic Price Competition With
An Application to Platform Design**

Andrew Sweeting
University of Maryland and NBER

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

A Model of Non-Stationary Dynamic Price Competition With An Application to Platform Design

Andrew Sweeting*

University of Maryland and NBER

September 2015

Abstract

I develop a tractable framework for conducting platform design counterfactuals in settings where many sellers compete and set prices dynamically, using approaches developed in the recent literature on Oblivious Equilibria. As an initial application, I use the model to study a simple platform design counterfactual using data from the secondary event ticket market on Stubhub.com where the perishability of the product being sold results in sellers facing a dynamic and non-stationary pricing problem. Currently, most transactions happen at low prices close to the event. Motivated by some simple theoretical examples, I investigate how the dynamics of prices, the timing of transactions, platform revenues and participant surplus would be affected if a commission structure that encouraged earlier transactions was introduced.

---

*Mailing address: Department of Economics, Tydings Hall, University of Maryland, College Park, MD 20742. Email: sweeting@econ.umd.edu. This project received funding from the NET Institute Summer Research Grant program. I am grateful to comments from several participants at the HEC Montreal summer IO conference, and especially Jean-Francois Houde who provided an excellent discussion. Several of these comments, and several different counterfactuals, will be included in future versions of the paper. All errors are my own.

1 Introduction

Many e-commerce platforms, such as eBay, facilitate trade between buyers and sellers without directly controlling prices, but collecting commissions which depend on the prices at which transactions take place. Platform owners have to make many choices over platform features - such as how to structure commissions and how to present search results to potential buyers - where the profits from alternative designs are likely to depend on how features affect equilibrium prices. For example, allowing potential buyers to identify a larger pool of listings that suit their preferences should increase transaction volumes, all else equal, but could potentially decrease equilibrium prices so much that platform revenues decrease. The same change would also tend to have ambiguous effects on the welfare of sellers, which might matter if the platform's success depends on encouraging more sellers to use it rather than a competitor.

In some settings, a standard Bertrand-Nash differentiated products model where sellers simultaneously set prices once-and-for-all might be sufficient to make reasonably accurate predictions about how equilibrium prices change. However, on many e-commerce platforms sellers face problems that are intrinsically dynamic in some sense that would significantly violate the assumptions of this model: for example, the products being sold (event tickets, accommodation, rental car reservations) may be perishable with sellers able to change prices, possibly many times, as the expiry date approaches; or alternatively, because they are small sellers who cannot constantly monitor prices they may change their price recognizing that the market may shift in some important ways before they change their prices again. In these settings, a change in platform design is likely to change the seller's value of being in the market in the future, so it would not be appropriate to treat the marginal cost implied by a price choice under an existing platform design as fixed and exogenous in the way that a static analysis typically would.

This paper sets out a model of non-stationary equilibrium pricing behavior that can be used to perform platform design counterfactuals in a setting where perishability gives rise to dynamic pricing incentives. The setting is the online resale market for tickets to sports events (specifically the current sample consists of a set of NFL games) on Stubhub, which is the largest platform for these transactions. For any single event (game, but I will refer to it as a 'fixture' to distinguish it from the pricing game between sellers that is the main subject of the paper), there are hundreds or thousands of sellers who list tickets in the months and weeks before the fixture takes place. As a fixture approaches, sellers cut prices systematically and by large amounts, and the frequency with which they change prices increases. This reflects the natural incentive of the seller of a perishable good to try to sell at a high price if he

can, but then to lower prices as his opportunities to try to sell their tickets in the future disappear. as demonstrated in Sweeting (2012), this logic leads to large pricing declines even in event tickets markets (in that case for Major League baseball games) where there are hundreds or thousands of sellers.

The key challenge in formulating a dynamic game is that computational needs, both in terms of CPU time and memory, grow rapidly in the number of players and states.¹ For example, as mentioned by Weintraub, Benkard, and Van Roy (2008), to store the value functions in a Pakes and McGuire (1994) Markov Perfect Nash equilibrium-type game with 20 symmetric players each with a state that can take on 40 possible values requires 20 million gigabytes of memory! Modeling a pricing game between all of the sellers on Stubhub for a particular fixture in this way would clearly be impossible, even if the game was stationary. In a non-stationary setting, necessary when products are perishable, the computational problem becomes even larger as it is necessary to solve for value functions in each period of the game.

To make the problem tractable, I follow ideas developed in the recent literature on Oblivious Equilibrium (OE). When OE is applied to a stationary game, the critical assumption is that, in equilibrium, each player maximizes its value given her own state, assuming that the rest of the market or industry is always in the average long-run state implied by equilibrium strategies. This reduces the size of the state space that has to be considered to the state space corresponding to the single-agent dynamic programming problem. In a non-stationary game, considered in the unpublished paper by Weintraub, Benkard, Jeziorski, and Van Roy (2008), the parallel assumption is that a player's strategy in any period t should maximize her value given her own state, assuming that the rest of the market or industry evolves deterministically along the (average) path implied by these strategies.² In return for these simplifications about players' expectations, and implicitly what they believe about their ability to influence how the market evolves, rich seller heterogeneity can be captured. For example, sellers can differ in the quality of tickets they have to sell, the number of listings that they have, and their expected values of being left with tickets at the time of the fixture which can also change over time. This is necessary to capture important features of the data in the current setting. This also opens up

---

¹As discussed in detail by Doraszelski and Pakes (2007) there are two problems created as the number of players grows. First, the rate at which the total size of the state space grows is much faster than linear. Specifically, even when we impose symmetry and exchangeability, the number of states is $\binom{N+k-1}{k}$ where N is the number of players and k is the number of states per player. Second, in a game where players states can change simultaneously, the number of future states that we can move to in the next period from any given current state will tend to increase the number of players. Therefore the computation time per state will also tend to increase.

²In the current setting, non-stationarity arises from the finite horizon nature of the game. In contrast, Weintraub, Benkard, Jeziorski, and Van Roy (2008) consider an infinite horizon game where there is a one-off shock, and the non-stationarity of the average state comes from the temporary transition to the new, long-run OE average state.

the possibility of considering platform design questions where different sellers are treated in different ways (for example, by giving discounts to sellers who have more tickets to sell, or by charging a lower percentage commission to sellers with higher quality tickets which will tend to sell at higher nominal prices), or where buyers are presented with search results that may differentially change the demand of different types of listing.³

In the current paper, the design change to be considered is an inter-temporal change to the commissions that Stubhub charges to sellers. While other design questions are also interesting and could well generate larger welfare or revenue effects, this counterfactual is natural to consider because it illustrates the advantages of having a computationally tractable dynamic model.⁴ During the period of the data, Stubhub charged a percentage commission to sellers that did not vary over time.⁵ In an environment where sellers set prices once-and-forever and had no value to being left with tickets at the end of the game, percentage commissions would have no effect on equilibrium prices. However, when sellers do potentially have some value to being left with tickets, a higher commission will be associated with a higher price, and, in a dynamic setting where sellers change their prices over time, sellers have an incentive to impose additional mark-ups the further they are from the event and these prices will be affected by the commissions that will be charged later. From an efficiency perspective, the fact that buyers who are in the market earlier on face higher prices than those who are in the market closer to the fixture creates an inefficiency, and, as I show in examples in Section 2, it is possible for the platform to raise the total combined surplus of buyers-and-sellers by changing the commission structure so that sellers set lower earlier prices, while not lowering platform profits. Alternatively, the platform can raise its revenues while not decreasing combined buyer-and-seller surplus.⁶ Therefore, the structure of the current paper is to estimate the parameters of the OE dynamic model using data where there is no variation in the commission structure, and then to re-solve the model allowing for intertemporal variation.

---

³In the current model I assume that buyers are aware of the quality of all listings, but, with richer data on search behavior, a more appropriate assumption would be to assume that buyers only consider a more limited subset of listings (consideration set) and that the design of the platform can affect which listings are likely to be included in the set.

⁴It is also practically relevant as the platform has recently experimented, on a limited scale, with some time-varying commissions. One objective is therefore to predict what might happen if this approach was rolled out as a platform-wide feature.

⁵For the majority of sellers the commission charged was 15%. Some larger sellers were charged lower commissions. Current estimation and counterfactuals are based on those that paid 15% commissions; however, all sellers will be included in the next iteration of the paper, and one can also look at counterfactuals that evaluate the effects of this type of cross-seller discrimination in commissions.

⁶As will also be illustrated below, changing the commission structure can also redistribute surplus from buyers-to-sellers or vice-versa. From a platform design perspective what is optimal will therefore also be impacted by whether it is more important to attract additional buyers or sellers to the platform, extensive margins that are not currently considered within the model. This may well depend on the exact nature of the event being considered.

The framework used in the current paper should be regarded as an initial investigation of the approach that shows promising results, with a model that could be extended in several directions. I plan to extend it several directions that will increase the framework’s novelty and its usefulness. Five directions stand out, and I will mention them here partly to highlight several limitations of the current framework that readers should be aware of.

First, as in the OE literature, sellers are assumed to believe that the state of the market, which here will be captured by the inclusive value of the available listings, evolves along a deterministic path, reflecting its ‘average’ path in equilibrium. This assumption is problematic in two respects: (a) the model predicts that the inclusive value should actually evolve stochastically, according to the stochastic arrival and purchase decisions of buyers, and the valuation shocks and price change opportunities received by sellers; and (b) there can be common shocks to the value of listings as news arrives about the performance of the teams evolved (do they have a chance to make the playoffs?, for example) and more prosaic factors such as the weather. The descriptive results here and in Sweeting (2012) both suggest that these factors actually have much smaller (and typically statistically insignificant) effects on prices than might be expected, but it would obviously be good to allow for some unexpected demand shocks within the model. A stochastic process could be handled using the ideas in recent work by Ifrach and Weintraub (2015), where players maximize against a chosen set of moments that describe the current and recent state of the market, rather than exact state that would capture the current prices and valuations of every seller.

Second, the effects of changing commissions on the rates of entry or exit decisions of both buyers and sellers into the market are ignored.⁷ This reflects the fact that there is little variation in the current data to identify these margins, but obviously these margins would have important impacts on the profitability of any changes (and could either reinforce or diminish the effects that are found here).

Third, I assume that buyers have logit demand over products, rather than (say) random coefficient demand where some buyers systematically prefer different types of tickets. Allowing for random coefficients of some kind (e.g., two types of consumers with systematically different preferences for tickets of different qualities) is feasible, although would increase the computational costs of solving the model.

Fourth, while the oblivious assumption seems a reasonable behavioral approximation for sellers with a single listing, it would be less appropriate for ticket brokers who list many tickets to the same

---

⁷This is an overstatement, as I do allow for sellers to delist tickets from Stubhub without being sold and this could be affected by the commission rate, although it will not be under the exact way that delisting is modeled at the moment.

event and may recognize that they can move the market. Indeed, Lee, Roberts, and Sweeting (2012) and Sweeting and Sweeney (2015) look at a set of events where a single seller accounts for almost 50% of transactions in the most desirable tiers of an arena, and around 30% of transactions for the arena as a whole, and provide some evidence that this seller can move the market and that this affects its pricing. For the events analyzed in the current paper there is no seller that is as dominant as this seller, but it would be interesting to extend the current framework to allow for at least some sellers to behave in a more sophisticated way. However, to do so would involve using a different framework to the approaches developed for applying OE to concentrated industries (Benkard, Jeziorski, and Weintraub (2013)) or Ifrach and Weintraub’s ‘moment-based Markov equilibrium’ approach, as these papers assume complete information so that all players are aware of the dominant players’ states. On Stubhub, in contrast, other sellers are not aware of which tickets are owned by larger sellers, and a new equilibrium concept would be required to deal with this informational environment.

Finally, I assume that buyers arrive stochastically and that, after making their decision about whether or what to buy, they exit the market forever. This assumption is typical in most of the revenue management literature, although recent work has allowed for strategic, long-lived buyers, typically in the context of monopoly (e.g., Hörner and Samuelson (2011)). I hope to relax this assumption with better data on the search behavior of individual buyers that should become available. Allowing for strategic buyers will have two effects on the welfare calculations. I will tend to overstate the benefits to lowering earlier prices if some buyers who do not buy then do buy when prices are lower close to the game, but, on the other hand, these buyers likely end up bearing significant search costs which it would be efficient to eliminate.

The structure of the current draft is as follows. The rest of this section discusses some related literature. Section 2 describes some examples that motivate the counterfactuals. Section 3 details the data and several stylized empirical facts that influence the structure of the model. Sections 4 and 5 present the model, and the solution and estimation strategies. Section 6 presents some initial estimates and Section 7 contains an illustrative counterfactual. Section 8 concludes.

The paper is related to several existing literatures. The first literature has been concerned with resale markets for event tickets. The central question in the theoretical literature (e.g., Courty (2003)) concerning these markets is whether it is efficient for secondary markets to exist at all, given that, historically, they have been subject to significant legal restrictions and hostility from primary market sellers, such as concert organizers and sports teams. Leslie and Sorensen (2014) estimate some of these effects, assuming a static model of how the secondary market clears, although allowing for those

people who buy in the primary market to change how much they value tickets before selling them in the secondary market. The current paper will consider the efficiency of how the secondary market clears in a dynamic setting, taking the existence of the market as given. The paper is also related to my earlier work in Sweeting (2012). That paper shows that individual sellers price dynamically and in a way consistent with simple models of revenue management (e.g., Gallego and Van Ryzin (1994), McAfee and Te Velde (2006)), using data on secondary markets for MLB tickets on Stubhub and eBay. It also provides evidence that the price that a particular seller sets has minimal effects on the prices that other sellers set subsequently, consistent with an OE type assumption, but it did not layout an equilibrium framework that could be used for counterfactuals. It also shows that, on eBay, data on buyers is consistent with the buyers who are currently active in the secondary market early on being people to live far away from the stadium and might want to make sure that they have tickets before making complementary investments.

The second literature relates to OE, and includes several papers by Weintraub, Benkard and co-authors laying out the concept in a variety of settings, and providing bounds that allow one to evaluate how much profit an individual firm might be able to gain if it unilaterally deviated from oblivious behavior to taking into account the true state of the industry. There have only been limited empirical applications of oblivious concepts to date (Xu (2008), Qi (2013), Saeedi (2014)). The most-closely related of these applications is Saeedi who uses a stationary OE framework to model the effects of making measures of seller quality available to buyers on eBay. One feature of these models is that it is assumed that players get to make moves in each period. In the current setting, players are observed to change prices infrequently (although the frequency increases as a fixture approaches). This will be interpreted as reflecting the fact that sellers only give the market limited attention and only occasionally consider changing their prices (as suggested by Ellison and Snyder (2014) as being the appropriate model for the sellers of memory chips on Pricewatch.com), and modeled as if sellers receive exogenous opportunities to move as in some of the literature on continuous time games (e.g., Arcidiacono, Bayer, Blevins, and Ellickson (2012)).

The third literature includes papers, including Saeedi's, that consider the design of platforms where buyers and sellers set prices. Dinerstein, Einav, Levin, and Sundaresan (2014) consider the effects of changing search protocols on eBay and show that a static pricing model, where buyers have consideration sets that are a function of the search results, can explain most of the observed changes in price levels. In the current setting a dynamic, and non-stationary, framework is required because of the perishability of the product being sold, although, as noted previously, I would also like to use the

current framework to consider what would happen if search options were changed. Hammond (2008) and Bauner (2015) also consider platform design on eBay focusing on the types of sales mechanism that sellers can use. Bauner’s paper is connected to the current paper in that it also considers markets for event tickets (using the same data as Sweeting (2012)), although it does not formally use a dynamic equilibrium concept. The particular type of counterfactual considered here, where I consider a change to the commission structure that may affect the path of prices over time, is also connected to the literature on dynamic/intertemporal price discrimination (e.g., Lazarev (2013)). In this literature, market structure is typically monopoly (as in an airline market served by only a single carrier), making the analysis more tractable. In the current paper the focus is on competition between many sellers, but where a platform may try to influence that competition.

2 Motivating Examples

In this paper, I estimate a dynamic equilibrium model in order to understand the possible benefits to adopting a time-varying commission structure. In this section I provide a very simple example to motivate this type of counterfactual.⁸

Suppose that there are two periods, $t = 1, 2$, so period 2 is the final period before the fixture. There are three (initially symmetric) sellers, each with one unit to sell, and one potential buyer in each period so if there is a sale in period 1, then there will be two sellers in period 2; otherwise there will be three sellers.⁹ A seller’s value to being left with an unsold unit at the end of the second period is either $v = v_2$ or $v = 0$ with equal probability. The realized value of v is revealed to the seller in period 2 but it is not known in period 1, capturing the fact that sellers may not yet know whether they can attend the game or whether they could sell the tickets to friends offline. From the perspective of the potential buyer in period t the listings are symmetrically differentiated and the utility from buying listing j is

$$u_{jt} = \gamma_t - \alpha_t p_{jt} + \varepsilon_{jt} \quad (1)$$

where ε_{jt} is a standard Type I extreme value (logit) error and p_{jt} is the listing’s price. In this example I assume that sellers simultaneously set prices each period, although this will not be true in the model that is taken to the data. As a baseline, I assume that the platform charges a commission of 25%

---

⁸A second motivation is that the platform has experimented with some time variation in commissions.

⁹You get the same qualitative results if the number of second period sellers is fixed at three irrespective of what happens in the first period.

of the sale price when a transaction occurs, and that there are no other fees.¹⁰ The counterfactuals in this example involve looking at how changes in the commission structure across periods affect the combined surplus of buyers and sellers, and the revenues of the platform.

Let us begin by assuming that $v_2 = 0$ (so that the objective of all sellers is revenue maximization throughout the game), $\gamma_1 = \gamma_2 = 3$ and $\alpha_1 = \alpha_2 = 4.5$ (i.e., demand is identical in each period). In this case, equilibrium prices in the first period are 0.47, and either 0.32 (no sale in the first period) or 0.39 (sale in the first period) in the second period.

Agent surplus (the sum of expected consumer surplus and seller revenues) is 1.50 (in money units) and expected platform revenue is 0.18. Note that the possibility of selling units in the second period causes sellers to charge a mark-up in the first period. This creates an inefficiency typically associated with third degree price discrimination where some buyers do not buy in the first period even though they would do so at equilibrium prices in the second period. It also hurts the platform because when setting this mark-up sellers fail to internalize the effect on the revenues of the platform. This mark-up and the associated inefficiency can be decreased if the platform charges a higher commission in the second period. When sellers are revenue-maximizing in the second period, a percentage commission has no effect on second period prices, but - by reducing the value of not selling in the first period - it reduces the prices that set in the first period.

In this simple example, a commission structure of 14% in the first period and 40% in the second period maximizes combined agent surplus holding platform revenues fixed, increasing surplus by 0.5%. On the other hand, a commission structure of 16% in the first period and 39% in the second period, increases platform revenues by 3%, while keeping surplus fixed at its baseline level. Note that: (i) the probability that a transaction happens in the first period is high in equilibrium (0.8774), limiting the possibility of increasing sales in the first period; and, (ii) the price decline is much smaller than observed in the data, with both factors suggesting that much larger welfare changes might be observed in the empirical example.

Other possible features of the market can also lead to this type of commission change having greater benefits. As an example, seller uncertainty about v . For example, if it is known in the first period that $v = 0.4$ (for sure), then there is no commission structure that can increase both agent surplus and platform revenues above their baseline levels simultaneously. However, suppose that in the first period sellers believe that $v = 0$ or 0.8 with equal probability (so that the expected v is still 0.4). Then in this case, a change in the commission structure to 18% in the first period and 40% in the

---

¹⁰In the data the platform charged a 10% commission to buyers, a 15% commission to most sellers and a lower commission to the largest sellers.

second period can increase platform revenues by 5.6% (holding agent surplus fixed) or increase surplus by 1% holding revenues fixed (commission structure 16% and 40%). This happens because, when v is uncertain, option value considerations tend to increase first period prices even further (in this case by just over 5% under the baseline commissions relative to the case when v is known to be 0.4 with certainty), and once again prices are increased without internalizing the lost surplus of buyers who would purchase at lower prices or the revenues of the platform.

Of course, lowering the first period commission and raising the second period commission can also redistribute surplus from sellers to buyers. Returning to the case where the sellers know $v = 0$, the seller’s expected surplus can be increased, holding platform revenues fixed, by raising the first period commission, and lowering the second period commission (this acts to raise first period prices, which, for standard reasons, competing sellers would like to do). For example, a commission structure of (35%, 10%) increases expected seller surplus by 5%, and a commission structure of (40%, 12%) increases platform revenues by 17%, while keeping sellers at their baseline welfare level.

3 Data and Motivating Stylized Facts

In section I introduce the data and emphasize several characteristics that motivate the model presented below, specifically: price dynamics and the concentration of sales and price changes at the end of the pricing game (close to the fixture); and, the importance of systematic seller/listing heterogeneity, which one could only hope to capture using something like an ‘oblivious equilibrium’ concept which simplifies the modeling of strategic interactions between sellers.

3.1 Data

The primary datasets are provided by Stubhub, the largest online secondary market for event tickets. Stubhub operates as a platform where sellers list tickets for seats in particular sections and rows at fixed prices, which they are able to change at any point in time. In what follows, a listing will be defined as a set of tickets sold by a seller (identified by an id number) in a particular section and row.¹¹ A feature of Stubhub - which makes it different from eBay - is that seller identities, or information

---

¹¹Note that this involved aggregating across different listing identification numbers in the data where a seller changes price (or another listing characteristic) by ending one listing and creating a new one, almost immediately. If there is a long gap (e.g., several days) between the old listing being removed and a new one being created then I treat the new listing as an entirely new listing. If a seller has two separate listings at the same time for the same fixture, section and row then I count them as a single listing if they have the same prices, and as different listings if they have different prices.

on seller experience or performance, are not revealed to buyers, with Stubhub providing a guarantee that tickets at least as good as those purchased will be given to the buyer if the seller fails to deliver, thereby removing the need for buyers to be concerned about seller reputation. In one sense, the fact that seller identities are not easily observed provides some justification for adopting an oblivious equilibrium-type of assumption as anonymity makes it difficult for sellers to keep track of how the states/prices of particular other sellers, including some who might be professional ticket brokers, are changing. Shipping fees are set by Stubhub and vary depending on the format of the ticket (electronic or paper), and, for paper tickets, the speed with which they are delivered.

The current draft makes use of data for 15 home regular-season games of two NFL teams in 2012, although most of the stylized facts also hold for all of the other (sports) events that I have looked at. The data comprises three datasets: (i) listings data; (ii) transactions data; and (iii) visit/sale conversion rate data. The transactions data provide information on every sale made on Stubhub: the exact time of the sale (to the second); ticket, seller and buyer id numbers, information about the listing - such as section, row and number of seats purchased; and, price and shipping information. The listings data tells us the characteristics of every available listing (e.g., game; section, row, number of tickets, information on whether a smaller number of tickets can be purchased (e.g., 2 seats from a five seat listing); price; shipping options) when the listing was initially posted and at any point where any of the characteristics of the listing changed. So, for example, there is a new record when a listing is first created; when any or all of the tickets are purchased; when any tickets are added or removed by the seller; or, when there is a price change. Once again, I know the time of the change down to the second. I can use this information to reconstruct exactly what listings were available at any point in time. Listings for seats in boxes/luxury suites and parking passes are excluded. I have added additional data on single-game face values in the primary market.¹² The visit/sale conversion rate data provides information on the number of visits and the proportion of visits that result in sales as a game approaches, although (in the current draft) these numbers come from a wide range of sports events rather than the specific fixtures used in the rest of the analysis.¹³ The key fact in this data is

---

¹²For the NFL teams that I examine, very few single game tickets are made available in the primary market in the weeks leading up to the game. One interesting direction to push the model would be examine interactions between the primary and secondary market in settings where the primary market is not sold out, especially in settings where the event organizer uses some type of dynamic pricing in the primary market.

¹³In the current data, I also do not know how many visits are by different potential buyers, rather than the same potential buyer visiting multiple times. The data is therefore interpreted as if the visits are by different potential buyers, but I hope to revisit this with better data at some point, which will allow us to get at the role of ‘strategic buyers’, who play an important role in the recent theoretical revenue management literature in both Operations Research and economics (e.g., Hörner and Samuelson (2011)). Note, however, that in these papers it is assumed that it is costless for buyers to continuously monitor the market. This is likely a far from reasonable assumption, and the evidence in Sweeting (2012) suggests that actually strategic timing behavior on the part of buyers may simply result in knowledgeable, flexible

that the probability that a visit results in a sale increases steadily from around 0.06 60 days before a fixture, to around 0.08 10 days before a fixture, before increasing to about 0.13 in the final day before the fixture.

3.2 Summary Statistics

Secondary markets for NFL tickets are large, so that there are many observations even though we are only using 15 games. In total there are 69,131 different listings (a seller-section-row combination), for tickets with combined face values over \$31.9 million, and 52,962 transactions where 43,606 different buyers pay a total of \$26.2 million (excluding shipping fees). The listings are posted by 10,293 different seller (ids) (and there are 7,375 different sellers in the transactions data) so that these markets are unconcentrated by the standards of most markets considered by IO economists, although there are some larger sellers, with one seller involved in 990 transactions for a single game (the average number of transactions per game is over 3,530). In the model, all listings will be modeled as having different sellers, although some parameters will be allowed to vary in order to capture the fact that large sellers do behave somewhat differently to small ones (in particular, they change prices more often and more likely to revenue-maximizing with little value of being left with tickets are game-time).

Table 1 contains some summary statistics on the data, based on transactions. Columns (2)-(5) are based on the tickets that are actually purchased, whereas columns (6)-(9) are based on all of the listings that appear in the buyer’s choice set, which is defined as including any listing from which the buyer could have purchased the same number of tickets that she actually purchased, taking into account any restrictions imposed by sellers on selling subsets of tickets from their listings.¹⁴

On average, a transaction takes place one month before a fixture. However, as illustrated in Figure 1, transaction times are skewed, so that one half of transactions happen within two weeks of the fixture, and around one quarter of transactions occur in the last three days. The model will be estimated focusing on the last 50 days, when almost 80% of all of the transactions occur. There are some differences in when transactions occur for different numbers of tickets: for example, 40% of purchases of two seats happen in the last week before the game, compared with only 27% of six ticket purchases. This difference will likely reflect in part the fact that a large group will likely have to make more complementary investments to attend an event (e.g., coordinating travel plans), which will tend

---

buyers only choosing to visit the market close to the event when prices tend to be low.

¹⁴So, for example, if we consider a transaction where 6 tickets were purchased, the choice set would include all of the listings with exactly six seats that were were available when the transaction took place, plus listings with more than six seats from which exactly six seats could have been purchased.

Table 1: Summary Statistics for 55,101 Transactions

Variable	Purchased Tickets				All Listings in Choice Set			
	Mean	Std. Dev.	Min.	Max.	Mean	Std. Dev.	Min.	Max.
Days Before Game	29.3	38.5	0.00	248.5	-	-	-	-
Size of Choice Set	-	-	-	-	744.8	624.8	1	2,566
<i>Ticket Characteristics</i>								
Middle Tier	0.14	0.35	0	1	0.17	0.38	0	1
Upper Tier	0.47	0.50	0	1	0.42	0.49	0	1
Front Row of Section	0.03	0.18	0	1	0.04	0.19	0	1
Row Number	16.0	9.95	1	48	15.5	9.9	1	48
Number of Seats Purchased	2.8	1.6	1	27	-	-	-	-
Number of Seats Available	-	-	-	-	4.5	4.0	1	64
<i>Price (\$ per seat, no shipping)</i>								
Face Value	129.61	93.70	50	700	140.55	100.47	50	700
Sale Price	180.42	145.35	1	2,400	-	-	-	-
Listing Price	-	-	-	-	234.37	186.67	1	2,500
<i>Shipping</i>								
Electronic download	0.63	0.48	0	1	0.50	0.50	0	1
Fedex/UPS	0.24	0.47	0	1	0.48	0.50	0	1
Last Minute	0.03	0.16	0	1	0.02	0.14	0	1

Figure 1: Timing of Transactions and Choice Set Size as a Fixture Approaches

to require more forward planning.

The size of the available choice sets also vary with the size of the purchase: on average, someone buying two seats has, on average, 1,022 listings to choose from, whereas someone buying six seats has only 248 options to consider.¹⁵ The size of choice sets also differs systematically over time, creating another source of dynamics and also suggesting that it might be possible to increase buyer surplus by encouraging more transactions in the weeks leading up to the event when there is greater product variety. Figure 1 also shows how the average size of the choice set varies with the time until the fixture for a buyer purchasing two seats. On average, the size of the choice set is 42% larger twenty days before the game than four days before the game. The decline in the size of the choice set reflects both the fact that as transactions occur the stock of available listings necessarily decreases, but also the fact that as a fixture approaches listings may be delisted without being sold. One reason for this is that listings have to be delisted when their shipping options do not provide enough time for them to be delivered to the seller (for example, no paper tickets being shipped by Fedex and UPS are not available within 3 days of the fixture, whereas tickets being provided for Last Minute delivery are available right up until the game), but many sellers also delist their tickets voluntarily before this point, which could reflect sales being made on other online platforms, offline sales or a decision on the part of the seller that they will attend the game themselves. Delisting will also be allowed for in the model below.

All of the games in the sample are played in stadia with three tiers, with most of the purchased and listed tickets being available in the upper and lower tiers. On average, prices are higher in the lower and middle tiers (average transaction price is \$243 per seat, compared with \$140 in the upper tier). The average size of a transaction is 2.8 seats, with purchases of two and four seats accounting for 57% and 17% of transactions respectively. Sale prices vary widely, with the average sale price (excluding shipping and the buyer's commission), being 40% above face value. Shipping options can be important to sellers, and also affect prices (electronic download is the cheapest at \$4.95, while Fedex can cost as much as \$25 per transaction depending on whether next day delivery is required). Some tickets are supplied to Stubhub by sellers in which case they can be sold up to game time and collected from a location close to the stadium. On average, and without controlling for ticket characteristics, sale prices are significantly lower than the average list price of a ticket in the buyer's choice set, while purchased tickets are much more likely to be delivered electronically, partly because they are more

---

¹⁵The large size of these choice sets would also make it relevant to think about the effects of search tools or other platform changes that might effect consumer demand. Currently the default sort order is lowest price first, which helps to explain why demand looks so elastic.

likely to be available close to the game.¹⁶ As we will see below, the estimated demand model, where ticket characteristics are controlled for, also shows that demand is highly price-elastic partly reflecting the fact that the default ranking on Stubhub, at least during this period of time, is ‘lowest price first’.

3.3 Price Dynamics

As mentioned in the Introduction, in these markets most transactions take place when prices are low close to the game. I illustrate the changes in prices using a regression framework

$$p_{ift} = X_{ift}\beta + FE_{if} + \varepsilon_{ift} \quad (2)$$

where the price variable is the price of a listing divided by face value (here, the price excludes shipping and the buyer’s commission), and the X variables contain listing characteristics (where applicable) together with dummies measuring the number of days until the game (0 to 1 days before the game excluded). The fixed effects vary according to the specification. The coefficient estimates are presented in Table 2.

The coefficients from columns (1) and (2) are based on regressions with fixture fixed effects, for either available listings (one observation per listing-day where we use the listing’s price at the beginning of the day) or transactions, where we control for observable ticket characteristics (number of seats purchased/available, section, row characteristics, tier etc.).

To control for team performance I include, for each team, dummies for whether the season is actually underway and then linear and quadratic terms that measure the number of weeks into the season. In none of the specifications below are these coefficients are significant. This may seem surprising, but in other analysis using secondary market data I also find at most very small effects on price of a variety of team performance measures (e.g. games back or games ahead in division, sports book odds of the team reaching the playoffs) in both MLB (Sweeting (2012)) or indoor sports (NBA/NHL) (Sweeting and Sweeney (2015)), once fixture fixed effects are controlled for.¹⁷

The time coefficients indicate that prices fall very significantly as a game approaches. For listings, average prices fall by 52% of face value in the 40 days before the fixture, with a 36% (of face) decline in the last two weeks. Transaction prices fall 44% in the last two weeks with statistically weak evidence

---

¹⁶There is a further distinction between electronic tickets that are available for instant download and ones that are not available for instant download. Instant download tickets can remain listed until closer to the game and may also be more attractive for consumers making last minute purchases.

¹⁷Of course, team performance could affect demand even if it has small effects on prices. Once hourly data on customer visit rates for each game is available this information it will be straightforward to test whether this is important.

of a small decline prior to that point. The characteristics coefficients (not reported) also indicate that there is lot of heterogeneity in prices even when the price is normalized by face value: for example, tickets in the front three rows of a section tend to sell for 20% of face value more than a fourth row seat, and 30% more than a 14th row seat. In column (3), I repeat the listings regression with game-section fixed effects, and find a similar price decline to column (1).

One can also look at how the prices of individual listings change, as the regressions in columns (1) and (3) partly reflect a composition effect as the set of listings that are available changes (in particular, cheap listings tend to be sold and so disappear from the sample, so that at least part of the composition effect tends to increase the price of available tickets as a game approaches). In this data, over 85% of observed price changes are price reductions. In column (4), I repeat the regression in column (3), but include seller-fixture-section-row (i.e., listing) fixed effects, so that the coefficients describe how the prices of listings that experience price changes, change. Here we see that sellers cut prices even more dramatically than the results in columns (1) and (3) suggest, with prices being lowered by more than 20% of face value from 40 days before the fixture to 20 days before the fixture, and falling by around 55% of face value in the last 20 days.

Obviously, these regressions measure mean price changes and one could imagine that the decline might be driven primarily by a collapse in the prices of very cheap tickets. As evidence on how the entire distribution of prices changes when a fixture approaches. Table 3 reports the results when we estimate the quantile regression equivalent of the model in column (1), for the 10th, 25th, 50th, 75th and 90th percentiles of the price distribution. In all cases prices fall significantly as a game approaches, indicating that it is not only very low prices that tend to fall.

One feature of the model presented below is that it allows for different sellers to have systematically different values of being left with listings unsold when the fixture takes place or they are no longer able to list their tickets because their shipping options would not allow delivery, so that sellers will tend to set a range of prices for similar tickets. To show that this is a feature of the data, I first calculate the residuals from the listing price regression in the first column of Table 2, which provides a measure of price that controls for listing characteristics. Figure 2 shows a plot that shows these residuals plus the relevant time effect (as a normalized price measure) for a random sample of individual listings to a particular game in a set of sections with face value \$155. As can be seen the typical pattern is that listings that are high/low in the price distribution tend to remain high/low in the price distribution even when they change prices. Close to the game, price decline predominate, whereas earlier on there are a mixture of price declines for some listings and price increases for others, suggesting that

Figure 2: Price Paths for a Sample of Listings with Face Value \$155 for One Game

heterogeneous and time-varying seller valuations of keeping hold of tickets may be important. These can be captured in the model below.

Note that a very small fraction of listings (less than 2%) change their prices automatically (an example can be seen at the top of Figure 2), as the seller selects an automatic, linear price decline, which reduces the price, usually every day, while the listing remains unsold. Excluding these listings does not change any of the basic results described above, and these listings are excluded from the analysis below, although it is perhaps an open question why this option is not used by more sellers given that the tendency to cut prices is so strong.

A very strong assumption in the model will be that sellers correctly predict how the market is going to evolve, and I estimate this predicted evolution by pooling fixtures (note that there will be a game fixed effect to account for some differences in levels across events). One simple test of how this is appropriate is to look at whether price-cutting behavior is similar across markets. I have therefore repeated the seller-fixture-section-row fixed effects regressions for each fixture separately, with more aggregated time dummies, specifically, 2-7 days to go, 8-16 days to go, 17-30 days to go and 31 plus. For all fixtures, the price declines are economically and statistically significant.¹⁸ The mean decline in from 2-7 days to 0-1 days is 36% of face value, with a std. deviation of 16%, and a range of 21% to 72%. The mean decline from 31 plus to 2-7 is 43% of face value with a std. deviation of 18%. Obviously these results suggest that cutting prices are systematic, but that there is also some heterogeneity across games. In later drafts, I would like to capture some of this heterogeneity using a more flexible assumption about how expectations of either aggregate demand or the state of the market evolve, following Ifrach and Weintraub (2015). However, one difficulty in doing this is that factors that one might have hoped would drive different declines - such as shocks to team performance - have remarkably little effect in this data.

4 Model

4.1 Overview

Consider the pricing game played by sellers for a particular fixture f . They play a discrete time game with a finite horizon. The time periods, $t = 1, \dots, T$, correspond to hours prior to the fixture, where T is the final period. Empirically I allow for different types of sellers, differing, for example, in their

---

¹⁸Statistical significance is based on standard errors that cluster on the section, as the previous form of clustering, on the fixture, is no longer available when the specification is estimated for each fixture separately.

Table 2: Descriptive Price Regressions: all standard errors are clustered on the fixture

Days to Fixture (0-1 excluded)	(1)	(2)	(3)	(4)
	Listings Fixture FEs	Transactions Fixture FEs	Listings Fixture-Section FEs	Listings Slr-Fix.-Section-Row FEs
2-4	0.123*** (0.025)	0.194*** (0.051)	0.135*** (0.026)	0.170*** (0.027)
5-7	0.243*** (0.033)	0.315*** (0.052)	0.260*** (0.037)	0.317*** (0.038)
8-10	0.330*** (0.036)	0.418*** (0.066)	0.351*** (0.041)	0.407*** (0.040)
11-13	0.362*** (0.049)	0.438*** (0.069)	0.384*** (0.054)	0.455*** (0.051)
14-16	0.371*** (0.055)	0.463*** (0.077)	0.394*** (0.060)	0.484*** (0.053)
17-19	0.388*** (0.059)	0.426*** (0.072)	0.412*** (0.063)	0.512*** (0.057)
20-22	0.401*** (0.061)	0.410*** (0.084)	0.427*** (0.065)	0.540*** (0.059)
23-26	0.421*** (0.060)	0.450*** (0.081)	0.447*** (0.063)	0.569*** (0.059)
27-30	0.438*** (0.058)	0.473*** (0.083)	0.464*** (0.060)	0.597*** (0.057)
31-35	0.447*** (0.057)	0.453*** (0.088)	0.474*** (0.059)	0.626*** (0.055)
36-40	0.444*** (0.059)	0.415*** (0.088)	0.472*** (0.060)	0.645*** (0.056)
41+	0.517*** (0.076)	0.453*** (0.095)	0.553*** (0.073)	0.763*** (0.066)
Other controls:	Number of seats, section, row, face value, shipping team-time/season over 3.1 m. listing-days	Number of seats, section, row, face value, shipping team-time/season 52,366	Number of seats, row, shipping options team-time/season over 3.1 m. listing-days	Number of seats, shipping options team-time/season over 3.1 m. listing-days
Obs.				

Table 3: Descriptive Price Regressions - Quantiles: NB. SEs are not clustered.

	(1)	(2)	(3)	(4)	(5)
Days to Fixture (0-1 Excluded)	10th Percentile	25th Percentile	50th Percentile	75th Percentile	90th Percentile
2-4	0.129*** (0.00575)	0.134*** (0.00547)	0.148*** (0.00647)	0.156*** (0.00981)	0.160*** (0.0137)
5-7	0.257*** (0.00609)	0.264*** (0.00501)	0.266*** (0.00701)	0.272*** (0.00979)	0.264*** (0.0140)
8-10	0.369*** (0.00635)	0.354*** (0.00528)	0.340*** (0.00653)	0.340*** (0.00869)	0.328*** (0.0116)
11-13	0.410*** (0.00628)	0.393*** (0.00549)	0.379*** (0.00556)	0.369*** (0.00966)	0.347*** (0.0147)
14-16	0.432*** (0.00629)	0.416*** (0.00513)	0.402*** (0.00536)	0.388*** (0.00834)	0.358*** (0.0127)
17-19	0.449*** (0.00518)	0.427*** (0.00462)	0.409*** (0.00585)	0.390*** (0.00727)	0.363*** (0.0111)
20-22	0.469*** (0.00705)	0.444*** (0.00560)	0.427*** (0.00471)	0.414*** (0.00833)	0.383*** (0.0137)
23-26	0.486*** (0.00554)	0.460*** (0.00438)	0.442*** (0.00548)	0.428*** (0.00863)	0.403*** (0.0126)
27-30	0.497*** (0.00673)	0.469*** (0.00496)	0.451*** (0.00561)	0.440*** (0.00771)	0.419*** (0.00989)
31-35	0.510*** (0.00591)	0.483*** (0.00537)	0.461*** (0.00631)	0.447*** (0.00776)	0.428*** (0.0131)
36-40	0.512*** (0.00628)	0.488*** (0.00479)	0.469*** (0.00571)	0.462*** (0.00790)	0.433*** (0.0128)
41+	0.517*** (0.00626)	0.498*** (0.00456)	0.483*** (0.00580)	0.472*** (0.00817)	0.446*** (0.0128)
Other controls:	Number of seats, section, row, face value, shipping, game, team-time effects				
Obs.	559,426 (50% sample of listing-days in the last 50 days)				

experience (a proxy for whether they are a ticket broker) and the shipping options of their tickets, and many of the parameters of the model can be allowed to differ across types. However, I will present the model as if there is only one type to economize on notation. There are a set of sellers, $\mathcal{S}_f$, who enter the market exogenously.¹⁹ Sellers can differ in the quality of their tickets (μ_j) and the shipping options associated with their tickets ($ship_j$), which are fixed attributes. The time-varying state variables of a seller are: the quantity of tickets (i.e., seats) they have left to sell (q_j); their perceived value of being left with tickets at the end of the game (v_j , which is a per ticket value); and their current price (p_j). The control/policy variable of each seller is its price. Sellers receive a payoff that is equal to their revenues from selling tickets (net of commissions and shipping); the value of tickets left unsold at the end of the game; and, a payoff they receive if their tickets are delisted without being sold.

4.2 Timing

Within each period, there are three types of events that occur in the following order, so one can think of each of these events as happening in its own sub-period.

1. Price change opportunities: incumbent sellers stochastically receive a price change opportunity with probability λ_t . When they receive a price change opportunity they also receive an updated signal on their value to being left with tickets v_j (as explained below, this allows us to explain any price that is chosen by a particular realization of v). The conditional cdf of v_{jt} is $F_t(v_{jt}|v_{jt-1})$ where v_{jt-1} is the value that v had in the previous period, but it may have been received many periods previously. F_t is strictly increasing in the sense that $F_t(v|v') > F_t(v|v'')$, for all $v, t, v'' > v'$, so that a higher perceived value today is more likely to be associated with a higher perceived value when a new price change opportunity arises. A new seller has an opportunity to set its price as soon as it enters.

2. Delistings without sale: tickets may be delisted because they are sold elsewhere or the seller decides to go to the game rather than selling it in the secondary market. Therefore delisting is likely to be a function of v_j . In the current draft I assume a particular model for delisting so that the possibility of being delisted has no direct effect on pricing behavior on Stubhub. Specifically, I assume that what happens is that someone who knows the seller's state approaches the seller and makes the seller a take-it-or-leave-it offer to buy all of the tickets in the seller's listing at a price equal to seller's valuation of keeping his tickets listed on Stubhub, at a price that the seller would be able to set immediately if he rejected the offer.²⁰ The probability that this event occurs is modeled using

---

¹⁹In practice we can only consider a limited number of time periods (50 days in this draft), so we assume that all sellers who arrived before this date, first arrive in period 1 and set prices in that period.

²⁰This formulation would make sense if we think of a price change opportunity as arising whenever a seller thinks

as a logit function, $H_t(\mu_j, q_j, v_j)$.

3. Consumer demand is realized. A potential buyer i who arrives on the site (specifically the buy page of the game), receives the following utility if she buys listing j

$$u_{ijt} = x_j\beta - \alpha_t \log p'_{jt} + ship_j\beta_t^{ship} + \varepsilon_{ijt} \quad (3)$$

where p'_{jt} includes the commission paid by buyers, and the *ship* variable captures the non-price effect of shipping options. If she buys nothing, she receives utility

$$u_{ijt} = \gamma_t + \gamma_f + \varepsilon_{i0t} \quad (4)$$

Many potential buyers may arrive during a period. The realized number of buyers is given by a Poisson distribution, with a time-varying parameter, and each buyer is interested in a particular number of tickets and only considers listings that allow him to buy this number of seats (which in practice will mean all listings that have more than the desired number).²¹ The inclusive value of the available listings for a given number of seats can be expressed as

$$I_{qt} = \log \left(\sum_{j \text{ has } q \text{ seats}} \exp(x_j\beta - \alpha_t \log p_{jt} + ship_j\beta_t^{ship}) \right) \quad (5)$$

and will play an important role in what follows.

4.3 Dynamic Pricing Problem of an Individual Seller

Define $V_{t+1}(q_t, v_t, p_t)$ as the seller's value at the beginning of the next period, where (q_t, v_t, p_t) are the inherited values of the state variables from period t . I now work backwards through period t to derive the seller's value, $V_t(q_{t-1}, v_{t-1}, p_{t-1})$. Let V_t^2, V_t^3 denote the seller's value immediately prior to the second and third sub-periods described in the previous sub-section. Note that the seller's value also depends on the expected path of the inclusive values, I_t , which the seller will assume he cannot affect, as well as the seller's time invariant characteristics.

3. Prior to demand. The probability that a buyer, who wants q seats and has listing j in her

---

about his Stubhub listing, as receiving an offer would likely trigger this type of consideration. However, the assumption is primarily for convenience.

²¹For reasons that will be explained below, I only consider buyers who are interested in buying an even number of tickets, and I also restrict attention to listings that initially have an even number of seats. I assume that it is always possible to buy a smaller even number of tickets from listings with an even number of tickets, which is typically true in the data. For example, over 90% of observed listings with 6 seats would allow people to buy 2 or 4 seats from the six.

choice set, would choose j will be assumed to be equal to²²

$$\frac{\exp(\underbrace{x_j\beta}_{\mu_j} - \alpha_t \log p_{jt} + ship_j \beta_t^{ship})}{\exp(x_j\beta - \alpha_t \log p_{jt} + ship_j \beta_t^{ship}) + \exp(\gamma_t + \gamma_f) + \exp(I_{qt})} \quad (6)$$

To define the value function, what is more useful is to define a function that gives the probability that the listing will q_t seats left at the end of the sub-period when it began with q_{t-1} , taking into account that many potential buyers, with different demands, may arrive. This will be denoted $G_t(q_t|p_t, q_{t-1}, I_t)$ where it should be understood that I_t is a vector. Then,

$$V_t^3(p_t, q_{t-1}, v_t) = \sum_{q_t} \{p_t''(p_t)(q_{t-1} - q_t) + V_{t+1}(q_t, v_t, p_t)\} G_t(q_t|p_t, q_{t-1}, I_t) \quad (7)$$

where p_t'' is the money, per ticket, that the seller actually receives when a sale is made at price p_t . In the final time period,

$$V_T^3(p_T, q_{T-1}, v_t) = \sum_{q_T} \{p_T(q_{T-1} - q_T) + v_T q_T\} G_T(q_T|p_T, q_{T-1}, I_T) \quad (8)$$

2. Prior to delisting. Given my simple model of delisting,

$$V_t^2(p_t, q_{t-1}, v_t) = H_t(\mu_j, q_{t-1}, v_t) V_t^3(p_t^*, q_{t-1}, v_t) + (1 - H_t(\mu_j, q_{t-1}, v_t)) V_t^3(p_t, q_{t-1}, v_t) \quad (9)$$

where p_t^* is the optimal price given the seller's state.

1. Prior to changing price. In the first sub-period a new realization of v_t and the possibility of changing price arises with time-varying probability λ_t . V_t is therefore

$$V_t(p_{t-1}, q_{t-1}, v_{t-1}) = \lambda_t \int \max_{p_t} \{V_t^2(p_t, q_{t-1}, v_t)\} f_t(v_t|v_{t-1}) dv_t + (1 - \lambda_t) V_t^2(p_{t-1}, q_{t-1}, v_{t-1}) \quad (10)$$

$$= \lambda_t \int \max_{p_t} \{V_t^3(p_t, q_{t-1}, v_t)\} f_t(v_t|v_{t-1}) dv_t + \dots \quad (11)$$

$$(1 - \lambda_t) H_t(\mu_j, q_{t-1}, v_t) \max_{p_t} \{V_t^3(p_t, q_{t-1}, v_t)\} + (1 - \lambda_t) (1 - H_t(\mu_j, q_{t-1}, v_t)) V_t^3(p_t, q_{t-1}, v_t)$$

Here we see that the assumptions about delisting mean that V^1 can be expressed in terms of V^3 terms that have already been calculated.

---

²²This is, of course, an approximation, as the inclusive value will itself be a function of the listing's own price. However, assuming this form greatly simplifies computation, and the quality of the approximation should increase as the number of other listings increases.

4.4 Non-Stationary Equilibrium

I now introduce the equilibrium concept. A seller's strategy consists of a pricing policy in each time period as a function of its state.

Definition 1 *Oblivious non-stationary equilibrium (ONSE).* An equilibrium comprises, (i) for each seller type, a pricing strategy that maximizes its value, in each state, given a commonly perceived path for $I_t, \hat{I}_t$; and (ii) a perceived path for $\hat{I}_t$ that is equal to the expected (i.e., average) path of I_t given the parameters of the model, the distribution of sellers' initial valuations v and quantities, and sellers' pricing strategies.

It is by using this concept that the model is tractable to both solve and estimate. However, even with this assumption, estimation much simpler when the seller's optimal price is a strictly increasing function of the unobserved state v as, in this case, I can, for a given guess of the parameters, invert the pricing functions to recover v exactly. Without this feature one would have to integrate out over the possible paths of v_t , for example using simulation.

While monotonicity of the optimal price in v seems an intuitive condition, similar to the condition that a firm's optimal price in a one-shot model is increasing in its marginal cost, it can fail in the dynamic model considered here when sellers have multiple units. To understand why this can happen, consider a simple example. Suppose that there are only two periods and that the seller, who begins with two seats, can choose its price in each period. Suppose also that in the first period, a significant number of potential buyers want a single ticket, whereas in the second period buyers are only interested in buying pairs. A seller with a high per-ticket v may be willing to sell a single ticket in the first period, whereas a seller with $v = 0$, who is trying to maximize her total revenues, may prefer to set a high price in the first period so that she knows she will have a product that she can sell in the second period.

When trying to estimate the model allowing for any integer number of tickets and any integer number of purchases, I found examples of non-monotonicity that seemed systematically to originate in this type of possibility. This problem partly reflects another limitation of the model: in reality, sellers can impose restrictions on whether a subset of tickets can be sold, but this feature is not allowed in the model as it introduces a combinatorial element that would be very difficult to handle. To move forward, I therefore restricted myself to considering inventories with even numbers of seats, and purchases involving even numbers of tickets, and verifying, as part of the computational algorithms for

solving and estimating the model, that monotonicity holds.²³ This is a simplification, but as around 80% of inventories begin as even numbers and around 80% of purchases are for even numbers, it does not do too much violence to the data. Formally showing conditions under which pricing functions are always monotonic under restrictions of this type is one of the next steps for this research.

5 Solving and Estimating the Model

In this section I describe how the model is solved (for counterfactuals) and estimated (to recover the parameters of the model, and initial distribution of v s for sellers). Estimation of some parameters requires the solution of an individual seller's dynamic programming problem.

5.1 Solution Method

The model is solved via an iterative process. Assume that I know the functions F_t and H_t ; the initial value of q_t and v_t for sellers arriving in the market (and their arrival time if they are not present in period 1); and the arrival rate of price/ v change opportunities. The iterative process is as follows:

1. Preliminaries.

- (a) Guess an initial timepath for each $\widehat{I_{t,q}^{iter=0}}$ where q is the number of available tickets in the listing;
- (b) Take a large number (S) of simulation draws for all of the stochastic events in the model (e.g., price changes opportunities for each seller);
- (c) Specify tolerances, (ϵ_I, ϵ_P) , that will be used to evaluate whether the model has converged; and,
- (d) Form an approximation to the G_t function for a range of different values of I_t for each t . This step is computationally intensive so it makes sense to only do it once and then interpolate to get values for a specific I_t . Recall that G_t reflects the probability that a seller sells a specific number of tickets (or pairs of tickets, given my restriction) given an initial q , μ and p . To calculate the exact probability of a particular quantity transition is not feasible (as many as 100 potential buyers may arrive in some hours), so instead I use simulation with the exception that I can calculate outcomes associated with no buyers or exactly one buyer entering the market analytically. I use 10,000 simulations (for each

---

²³Looking at purchases of $\{2,4,6\}$ seats, a convenient pattern in the data is that as a fixture approaches the proportion of two seat purchases tends to increase, as does the proportion of four seat purchases relative to the proportion of six seat purchases. This reduces the problem that a low v seller might want to high to avoid being left with an unsellable smaller number of seats.

hour) of the arriving customers and their purchase decisions for cases where more than two buyers enter the market, and calculated simulated probabilities of sale for the (μ, q, p, I, t) grid. The same draws are used for all (μ, q, p, I) for a given t in order to ensure that the probability that at least a given quantity of tickets is sold is increasing in μ and decreasing in p .

Note that in the current solution procedure, the majority of the computational time is taken in this step, both because of the size of the grid and the number of simulations that have to be taken. The complexity of this simulation is also one reason why I do not try to include strategic buyers, who could optimal time their purchases, into the model. With strategic buyers, sale probabilities would be a function of the entire time path of I s, rather than just the contemporaneous I_t s and computing sale probabilities with much higher dimensions is computationally infeasible.

Then **on iteration number k ,**

(a) given the path $\widehat{I_t^{iter=k}}$, solve for the dynamic programming problem for each type of seller to get the optimal pricing strategies. The problem is solved - for each type and for each period - on a grid of values (μ, v, q, p) . I use 21 grid points for μ , q can take on values $\{2, 4, 6, 8, 10\}$, and I use 57 points for v (ranging from -\$200 to \$1000) and 83 points for p (ranging from \$1 to \$2000, with uneven spacing so that there are more points up to \$800 which is where I observe most of the prices).²⁴ v and p are continuous, so I integrate out over possible values of v in the next period using the trapezium method, and then find the optimal price where the value function associated with different prices is approximated using a cubic spline. The resulting policy function is $p^{*, iter=k}(\mu, v, q)$.

(b) simulation to calculate a new path $\widehat{I_t^{iter=k}}'$. For each set of simulation draws, I can simulate what happens using the simulation draws, the optimal policy functions, and the G_t function to simulate a history of each listing. These paths can then be used to calculate a path of I for each simulation draw, and $\widehat{I_t^{iter=k}}'$ is simply the average across simulation draws.

(c) if $\max \frac{|\widehat{I_t^{iter=k}}' - \widehat{I_t^{iter=k}}|}{1 + |\widehat{I_t^{iter=k}}|} < \epsilon_I$ and $\max \frac{|p^{*, iter=k}(\mu, v, q) - \widehat{p^{*, iter=k}}(\mu, v, q)|}{1 + |p^{*, iter=k}(\mu, v, q)|} < \epsilon_P$, stop. Otherwise update $\widehat{I_t^{iter=k+1}} = \widehat{I_t^{iter=k}} + \frac{\widehat{I_t^{iter=k}}' - \widehat{I_t^{iter=k}}}{\gamma_1(k^{\gamma_2}) + \gamma_3}$ and continue to the $k + 1^{th}$ iteration. Currently $\gamma_1 = 5, \gamma_2 = 0$ and $\gamma_3 = 0$.

---

²⁴Given free disposal, one might disagree with allowing negative values for v . However, in the data some listings, with prices below the level that would maximize expected revenues, can only be rationalized with negative v s, and, at least a long time before the game, a negative v might be interpreted as being a revenue-maximizing seller who has some cost to maintaining his on Stubhub so is willing to accept a lower price.

5.2 Estimation

Estimation proceeds in several steps where I aim to recover as many parameters as possible without solving the dynamic programming problem. When we do have to solve the DP, the assumptions lead to a relatively convenient expression for the likelihood. Joint estimation would certainly be more efficient, and some joint estimation may be required with richer formulations of the model.

The steps are as follows.

5.2.1 Step 1: Buyer Demand, Inclusive Values and Buyer Arrival Rate

The data consists of very detailed (second-by-second) data on what was available each second to a buyer who wants a particular number of seats, details on exactly when transactions occurred and what was purchased, and much coarser estimates (based on a wider range of sports events in the same year as the events used in estimation) of the probability that a visit to the Stubhub buy page for the event resulted in a sale as a function of the number of hours until the fixture. Given that the different nature of the demand data that can be used, it makes sense to further separate estimation into stages. In the first step transactions and their corresponding choice sets are used to estimate the parameters of the utility function

$$u_{ijt} = x_j\beta - \alpha_t \log p_{jt} + ship_j\beta_t^{ship} + \varepsilon_{ijt} \quad (12)$$

using a standard maximum likelihood estimator, where we pool all of the NFL games in our sample. Here we use the fact that for a multinomial logit, the parameters can be consistently estimated using a subset of choices. Here we are using observations where a transaction actually took place, so that the ‘outside good’ is excluded.

I can then use these estimates to calculate the average value of I_{qt} , for all of the fixture-hours in our sample for $q = 2, 4, 6$. In practice, these values show similar intertemporal patterns (see below), but they differ somewhat in their levels across fixtures, presumably because games differ in their innate attractiveness. Therefore, the next step is to calculate fixture-time- q specific expected values of these inclusive values by regressing the observed values on a polynomial in time and game fixed effects (the coefficient on these will be the game attractiveness variable mentioned below). I run separate regressions for each q , and the expected values are simply the predicted values from these regressions.

The final step jointly estimates the Poisson parameter for expected number of buyers arriving in a particular hour; the probability that a buyer wants 2, 4 or 6 seats; and the parameters describing how the valuation of the outside good evolves. All of these probabilities are allowed to be a function of

the time until the event. The model gives an expected value for the number of purchases of a given quantity for a given fixture f : for example, for two seats,

$$\mathbb{E}(\# \text{ of two seats purchases}) = \underbrace{\lambda_{t,f}^{\text{buyer}}}_{\substack{\text{expected \# of} \\ \text{buyers visiting} \\ \text{SH in hour } t}} (\tau) \underbrace{\frac{\exp(X_{t,q}\theta)}{\sum_{q'=2,4,6} \exp(X_{t,q'}\theta)}}_{\substack{\text{prob. an arriving buyers} \\ \text{wants two seats}}} \underbrace{\frac{\exp(\widehat{I_{t,f,2}})}{\exp(\widehat{I_{t,f,2}}) + \exp(\gamma_t + \gamma_f)}}_{\substack{\text{prob. someone interested} \\ \text{in two seats make a purchase} \\ \text{rather than outside goof}}} \quad (13)$$

and we estimate the Greek-letter parameters by matching the observed number of purchases of each q in each hour with their expectations, as well as matching to probability that a visit results in a transaction to the conversion data.

5.2.2 Step 2: Arrival of Price Change Opportunities

A feature of the model - where pricing strategies vary with t , each price change opportunity is associated with a change in v , there are no costs to adjusting prices and prices are a continuous choice - is that, with probability one, a seller should change his price whenever he has an opportunity to do so. I can, therefore, simply estimate the probability that a seller gets to change her price from the observed frequency of price changes, and I estimate a simple logit (as polynomial in the number of hours until the fixture) for the probability that a price change opportunity arrives.²⁵

5.2.3 Step 3: Evolution of v , Delisting Probabilities and Initial Conditions

The most complicated part of estimation is the third step, because it requires solving the dynamic programming problem of a seller, although, unlike when solving the model for counterfactuals, I_t is treated as fixed along its estimated expected path. The grid used and solution method is the same as described in iteration step (a) of section 5.1, with a prior step being approximation of the G_t function.²⁶ The end date depends on the type of the listing: for example, for Fedex listings the end point is four days prior to the game.²⁷

²⁵The price change frequency can be made a function of μ and q without complication. If the price change frequency is made a function of v then step 2 must be integrated with step 3.

²⁶To solve the counterfactuals we have to calculate G_t for a variety of different values of I_{qt} for each qt (as the path of I will change during solution). For estimation we only have to approximate G_t for the expected path calculated in previous stages.

²⁷Currently I group listings based on the commission rate by the seller (which is also a proxy for the seller's size as it is ticket brokers who pay lower rates) and the shipping type (Fedex vs. electronic/LMS). One could make further distinctions - for example, between listings with overnight Fedex vs. 3-Day Fedex options - but currently I do not do

The parameters to estimate are those of the F_t function and the parameters of the H_t function that describe the probability of delisting. For given parameters, the model can be solved to generate an optimal pricing function in any period as a function of (μ, q, v) . If the pricing function is monotonic in v , then for a given (μ, q, p) , which is observed for each listing that changes price, we can compute the v associated with the chosen price by inverting the policy function. I form the following likelihood²⁸

$$\log L(\theta^v, \theta^{delist}) = \sum_{i=1}^{N_{obs}} \sum_{t=t_i^{first}}^{t_i^{last}} \left(\mathcal{I}(t > t_i^{first}) \left[\begin{aligned} &\mathcal{I}(i \text{ changes price in } t) \log \left(\lambda_t f_t(\widehat{v_{it}(p_{it})} | \widehat{v_{it-1}}, \theta^v) \right) + \dots \\ &\mathcal{I}(i \text{ no price change in } t) \log(1 - \lambda_t) \\ &\mathcal{I}(i \text{ no price change in } t) \log(1 - \lambda_t) + \dots \\ &\mathcal{I}(i \text{ delist in } t) \log(H(\mu_i, q_{it-1}, \widehat{v_{it}}, \theta^{delist})) + \dots \\ &\mathcal{I}(i \text{ no delist in } t) \log(1 - H(\mu_i, q_{it-1}, \widehat{v_{it}}, \theta^{delist})) + \dots \\ &\log(G(\mu_i, p_{it}, q_{it}, q_{it-1})) \end{aligned} \right] + \dots \right) \quad (14)$$

which is maximized using the simplex algorithm.

6 Estimates

I now present an initial set of coefficient estimates. I note where there are some simplifications from the methods laid out above.

6.1 Step 1: Buyer Demand, Inclusive Values and Buyer Arrival Rate

The estimates of the utility function coefficients are shown in Table 4. They are based on transactions for no more than 6 seats, and the price coefficient is allowed to vary with the number of days until the fixture.

The results indicate that demand is highly price-elastic (as the probability that any listing is sold is small given the size of the choice sets, the log(price) coefficients give a very close approximation to the price elasticity for a given listing). The absolute value of the elasticity is highest 7 days before the game. The price of shipping is included in the price measure²⁹, so the shipping coefficients should be interpreted as measures of convenience. The patterns are broadly intuitive in that instant download

---

so. With a four day endpoint I simply ignore what happens to Fedex listings that were able to remain longer.

²⁸Of course, for estimation only the first, third and fourth lines contain parameters to be estimated. However I include the other lines for completeness and for showing what the likelihood would be, for example, if the arrival of a price change opportunity was a function of unobserved v .

²⁹Of course, this assumes that buyers correctly account for the cost of shipping when making their choices, whereas the time of the data they only saw the cost of shipping added to the price once they had initiated the checkout process.

Table 4: Demand Estimates: Conditional Multinomial Logit Model

<i>Price</i>	<i>Row Variables</i>		<i>Face Value cont.</i>
Log Price (incl. shipping)	-6.397 (0.045)	Row Number -0.012 (0.0006)	\$155 5.072 (0.062)
* 0-2 Days Prior to Fixture	0.282 (0.048)	* Middle Tier -0.070 (0.002)	\$205 3.335 (0.076)
* 3-7 Days Prior to Fixture	0.429 (0.048)	* Upper Tier -0.035 (0.001)	\$255 2.249 (0.063)
I(more than 7 days prior)*	0.170	First Three Rows of Section	\$295
log(days prior - 7)	(0.012)		1.424 (0.031)
<i>Delivery (non-instant electronic excluded)</i>			(0.032)
Instant	-0.114 (0.018)	* Middle Tier -1.332 (0.038)	\$395 5.285 (0.077)
Fedex Delivery	-0.237 (0.015)	* Upper Tier -2.219 (0.036)	\$420 5.837 (0.066)
Last Minute (at stadium) Delivery	-0.866 (0.076)	<i>Face Value Group (captures tier)</i>	\$700
Instant * 0-2 Days Prior	0.572 (0.069)	\$50 -1.118 (0.056)	
Instant * 3-7 Days Prior	0.210 (0.033)	\$75 -0.829 (0.054)	
Fedex * 3-7 days	0.167 (0.038)	\$95 -0.727 (0.056)	
LMS * 0-2 Days	0.526 (0.103)	\$105 0.313 (0.054)	
		\$125 1.634 (0.051)	
No. of listings in choice sets: 39,236,521; no. of transactions: 51,465; log-likelihood: -277,642.0, pseudo-r ² : 0.133			

and being able to pick up tickets at the stadium are more attractive when the fixture is close at hand. Seats further back from the front of the section are less attractive, although it is not clear why seats in the front three rows are unattractive in the upper tier. The face value coefficients generally indicate that tickets with higher primary market face values are more valuable to buyers in the secondary market.³⁰

One comment is in order about the specification. All listing characteristics, including price, are treated as exogenous. Sweeting (2012), using data on MLB tickets on eBay, found it necessary to instrument for prices - using characteristics of sellers that correlate with the level of prices that they set, but which buyers should not care about - to find elastic demand. I will try a similar approach going forward, but I note, for now, that estimated demand in Table ?? is much more elastic than I found even with instruments in Sweeting (2012) and that it is much easier to control for characteristics that are observed by potential buyers on Stubhub than on eBay where many sellers write significant amounts of text describing their listings.

These demand estimates are then used to calculate the inclusive values, I_{qt} , for transaction quantities $q = \{2, 4, 6\}$. Time periods are hours, and we use the last 50 days prior to the game. Figures 3 and 4 show the observed time paths and the smoothed predicted values (used in the next stages) for $I_{2,t}$ and $I_{4,t}$ for several games in the sample. In general the match is good (and it is actually better for $q = 6$ where the observed paths of the inclusive value are quite flat), although a number of fixtures follow a path like game E in the $q = 2$ case where the inclusive value turns down in the last few hours before the fixture (associated with a smaller number of available listings, or higher prices), rather than up, which is the more common direction of change.³¹

The predicted inclusive value paths are used in the next step of demand estimation, where I estimate the parameters that reflect the mean utility of the outside option of not purchasing and expected number of potential buyers arriving on the site in each hour, together with the parameters of a logit function that determines whether the potential buyer wants 2, 4 or 6 seats. The parameter estimates (with standard errors clustered on the game using a bootstrap, but note that they do not correct for the estimation error in the first stage of estimation) are shown in Table 5.

While the estimates based on smoothed I are used in subsequent states, I also report estimates

---

³⁰Note that buyers do not see face values on Stubhub, and they are not necessarily easy for someone to find out when the primary market is sold out. Therefore I interpret these coefficients as reflecting ticket quality, rather than direct valuations of the face value of tickets.

³¹Discussions with Stubhub employees indicate that this variability in what happens in the last day or so is quite common across both sports and non-sports events, whereas what happens in the longer-run up to the event is generally seen as being much more predictable.

Figure 3: Inclusive Value Path for Two Seat Purchases

Figure 4: Inclusive Value Path for Four Seat Purchases

Table 5: Demand Estimation: Arrival Rate and Valuation of Outside Good. SEs clustered on the game

	<i>Arrival Rate Parameters</i>		<i>Outside Good Parameters</i>		<i>Number of Seats Buyer Interested In ...</i>	
	Smoothed I	Observed I	Smoothed I	Observed I	Smoothed I	Observed I
Constant	2.342 (0.124)	2.364 (0.175)	-18.956 (0.734)	-18.998 (0.995)	Constant for Buyer Wanting 4 Tickets	-0.864 (0.131)
Following variables are normalized to have mean zero, std. dev 1					* Hours/4	0.137 0.000
Hour	-0.148 (0.068)	-0.125 (0.150)	0.791 (0.231)	0.659 (0.178)		(0.075) (0.353)
Log(hour)	-0.122 (0.198)	-0.153 (0.261)	-0.638 (0.256)	-0.791 (0.211)	Constant for Buyer Wanting 6 Tickets	-2.013 (0.156)
Hour ²	-0.145 (0.089)	-0.126 (0.132)	0.727 (0.235)	0.656 (0.249)	* Hours/4	(0.227) -0.058
Hour ³	-0.109 (0.160)	-0.097 (0.165)	0.291 (0.219)	0.265 (0.196)		(0.203)
Hour ⁴	-0.079 (0.211)	-0.072 (0.183)	-0.121 (0.262)	-0.112 (0.091)		
Hour ⁵	-0.0499 (0.244)	-0.044 (0.194)	-0.449 (0.302)	-0.412 (0.189)		
Hour ^{$\frac{1}{2}$}	-0.121 (0.084)	-0.110 (0.197)	0.314 (0.108)	0.154 (0.128)		
Hour ^{$\frac{1}{4}$}	-0.131 (0.094)	-0.138 (0.182)	-0.112 (0.103)	-0.274 (0.171)		
Game Attractiveness (FE from I estimation)	0.354 (0.115)	0.376 (0.099)	0.612 (0.532)	0.518 (0.777)		

where I use the observed I for each game to provide a comparison, and the point estimates are generally very similar and individual coefficients are not significantly different from each other. The results are broadly sensible with a higher probability that buyers want two seats, especially close to the game, although the standard errors are large, which appears to reflect the fact that too many polynomial terms may, currently, be included.. Note that the game attractiveness variable, which comes from the game fixed effect in estimation of I_2 , can be high either when there are many tickets available, as tends to happen for the most in-demand games, especially several weeks before the event, or prices are low, so it is not a priori clear what its sign should be. Figure 5 shows the implied (based on the smoothed estimates) expected paths of the mean utility of the outside good and the expected number of potential buyers wanting either two or four seats who should arrive in a given hour for one of the games in the data.

6.2 Step 2: Arrival of Price Change Opportunities and Demand

As noted above, the probability that a seller receives a price change opportunity is estimated using observed price changes. At this point, I divide the data on sellers into different groups - by commission rate and delivery options - as it is plausible that they will have different pricing behaviors. The current paper reports estimates using data from game A. The data is organized so that for each listing (recall this is a seller-section-row-fixture combination) I see, in each hour, whether there was a price change, whether the listing was delisted and whether there were sales. A price change is based on the difference between the price at the beginning of the hour and the end of the hour even if several price changes occur (which is rare, except that sometimes a price change is immediately reversed), and, to be consistent with the model, all sales are assumed to be made at the end of hour price. Listings where only a subset of tickets are delisted are dropped.³² This gives an estimation sample of 746 listings, of which 160 are already present 50 days before the game.

The probability of price change is then estimated as simple logit function - with covariates, a constant, hours before fixture and $\log(\text{hours before fixture})$. The implied probability that the seller consider changes price in an hour (and so also receives a new draw of v) is 0.0006 50 days before the fixture, 0.01 ten days before the fixture, 0.047 four days before the fixture and 0.095 in the final day when the tickets must be delisted if they are not sold.

Before estimating the process for v , I also construct the G_t function which gives the probability

---

³²I also currently drop listings where a sale and a delisting takes place in the same hour as the current timing assumes that delisting would take place before a sale.

Figure 5: Expected Number of Buyers Per Hour and Valuation of Outside Good for Game A

that different quantity changes happen through sales. Figure 6 shows the probability that a listing with 4 seats goes to have having 4, 2 or zero tickets left, for either $\mu = 0$ (a moderately low value of quality, the lowest is -2) or $\mu = 8$ (a fairly high value, the highest is 10), 100 hours before the end of the pricing game (approximately one week before the fixture) or 100 hours from the start of the pricing game (approximately 43 days before the fixture). For lower prices, the probability of a sale goes up, and the graphs (where $\log(p)$ is on the y-axis) make clear how elastic demand is. When there are many potential buyers in the market (as in the $t = T - 100$ case), the probability that a listing would go from having 4 to exactly 2 seats over the hour is very small except for a very narrow range of prices: at a high price, no tickets will be sold, whereas at a lower price it is almost certain that all seats will be sold (either as two 2-seat transactions, or one 4-seat transaction). On the other hand, further from the game, when fewer buyers arrive, the range of prices for which one listing is likely to go from having four seats at the beginning of an hour to only having two is greater.

6.3 Step 3: Evolution of v , Delisting Probabilities and Initial Conditions

The estimates of the the process governing the evolution of v , and the logit parameters for the probability that tickets are delisted are as follows, where v is in tens of dollars. The results reported here are based on small sellers who ship their listings by Fedex:

$$v_{jt} = \underset{(0.0005)}{0.941} v_{jt-1} + \underset{(0.0363)}{2.414} + \underset{(0.0432)}{0.001} \frac{t}{T} + \underset{(0.0012)}{0.720} \mu_j + \varepsilon_{jt} \quad (15)$$

$$\text{where } \varepsilon_{jt} \sim N(0, \sigma_{jt}^2), \sigma_{jt} = \exp \left(\underset{(0.005)}{2.400} - \underset{(0.0005)}{0.00010} \frac{t}{T} + \underset{(0.001)}{0.031} \mu_j \right) \quad (16)$$

except that the v s are truncated to lie on the interval $[-\$400, \$1000]$, and

$$\Pr(\text{delist}_{jt} = 1) = \frac{\exp \left(\underset{(0.624)}{-7.774} - \underset{(0.004)}{0.002} v - \underset{(0.035)}{0.023} \mu_j - \underset{(0.058)}{0.029} q_{jt} - \underset{(2.4046)}{7.319} \frac{t}{T} + \underset{(1.548)}{10.088} \left(\frac{t}{T} \right)^2 \right)}{1 + \exp \left(\underset{(0.624)}{-7.774} - \underset{(0.004)}{0.002} v - \underset{(0.035)}{0.023} \mu_j - \underset{(0.058)}{0.029} q_{jt} - \underset{(2.4046)}{7.319} \frac{t}{T} + \underset{(1.548)}{10.088} \left(\frac{t}{T} \right)^2 \right)}$$

where standard errors are in parentheses.

Several aspects of these results are sensible: the v s are estimated to be highly serially correlated, which can explain why a seller who has a relatively high price will tend to remain a relatively high-priced seller if he changes his price; dollar values of v are increasing in the quality of the ticket (μ_j); the delisting probability in a period is small, but it increases quite rapidly as the end of the time

Figure 6: G_t Function: Red Line $\mu = 8$, Black Line $\mu = 0$

horizon approaches (the minimum implied by the quadratic in time is at 32 days before the fixture). v and ticket quality have no significant effects on the delisting probability, while the q_{jt} effect implies that sellers with a higher quantity of tickets remaining are less likely to delist. The least satisfactory aspect of the results may be that the standard deviation of ε_{jt} is estimated to be unrealistically large (e.g., $\exp(2.4) \approx 11$, implying a standard deviation of the innovations in v round \$110 dollars for a listing with quality $\mu = 0$). This reflects some listings have a new price that is so high [low] that the probability that they would sell in any subsequent period (without a price change, given the assumed path of I) is essentially zero [so high that sale is almost certain, and would also be almost certain at a higher, more profitable, price] implying that v must be high [strongly negative], even if the new price is only slightly higher than the previous price.³³

One feature of the results requiring further consideration is that the implied vs tend to increase over time. The reason for this is that the large value of σ implies that option value effects should have a large impact on pricing decisions, so that, for a constant v , sellers would, on average, tend to cut prices even more than they do in the data. The model explains why they do not by having low initial values of v . Looking at the implied vs , the values close to the game are almost all positive, whereas at the beginning of the pricing game the majority are negative (overall, the average implied v is -\$10). While the current estimates can be used to perform counterfactuals it may make sense to impose that, on average, sellers' beliefs about their vs are correct on average during estimation.³⁴

7 Counterfactual

As an application, I re-solve the model to examine how the timing of purchases, the equilibrium dynamics of prices and platform revenues would change if, motivated by the examples in Section 2, Stubhub changed its commission structure to encourage earlier trade. Obviously many changes are possible, and, rather than trying to search for the changes that would be revenue or surplus optimal, here I will focus on one example change.

At the time of the data, small sellers were charged a 15% commission whenever their tickets were sold. I therefore compute what would be expected to happen under an alternative regime where the commission is 7.5% until seven days before the event, and then increases 4.1 percentage points

---

³³Of course, the same logic holds in reverse. If the price subsequently drops to a level where a sale is quite likely - and as Figure 6 showed, this might not require too large of a drop - then it will appear that v has dropped quite dramatically, implying another large innovation.

³⁴It may also make sense to use a model where some sellers are very confident about their vs while others are very uncertain. Currently the parameters are estimated assuming that they are the same across all high commission-Fedex sellers.

today until the game itself. For the observed transactions in the data, this change would be essentially revenue-neutral for Stubhub, reflecting the fact that most transactions take place relatively close to the day that the fixture takes place. Following the solution recipe outlined in Section 5, I solve the model both for the constant 15% commission and for the alternative regime, and then in each case simulate outcomes. To do so, I assume that the market consists of 4,600 listings sold by small sellers with Fedex tickets (i.e., their parameters reflect those reported above), where 4,600 matches the average number of listings observed per game in the NFL data.

The results are reported in Table 6, where I report the results on a subset of days as the game approaches.

Table 6: Counterfactual: Results for Constant Commission (normal font) and Time-Varying Commission (italics)

	Days Before Fixture							
	40	30	20	15	10	5	3	1
Avg. List	225	226	220	214	208	196	175	160
Price (\$)	<i>214</i>	<i>214</i>	<i>210</i>	<i>206</i>	<i>190</i>	<i>186</i>	<i>180</i>	<i>169</i>
Pr(sale)	0.06	0.06	0.07	0.08	0.10	0.14	0.16	0.18
	<i>0.08</i>	<i>0.08</i>	<i>0.09</i>	<i>0.09</i>	<i>0.13</i>	<i>0.12</i>	<i>0.16</i>	<i>0.17</i>
Platform	1,350	2,600	4,160	2,590	3,130	7,740	10,520	14,350
Revenue (\$)	<i>920</i>	<i>1,780</i>	<i>2,740</i>	<i>1,700</i>	<i>2,100</i>	<i>6,910</i>	<i>16,840</i>	<i>33,480</i>
Δ C.S.	+3%	+1%	+4%	+2%	+4%	0.5%	-1%	-2%

As expected, the change in commission structure reduces how much prices fall over time. More than one week before the game, sellers are keener to sell which causes prices to fall, and the average probability that a listing is sold to rise, as some buyers (whose arrival is held to be the same as in the base case) choose to make a purchase rather than choosing the outside good. Platform revenues, however, tend to fall as the effect of the lower commission rate, combined with lower list prices, dominates the increased number of sales. On the other hand, close to the game prices are higher under the counterfactual as, at least for positive values of v , sellers become more willing to keep hold of their tickets rather than making a sale when the commission rate is high (the maximum is over 35%). On the other hand, because the vs of sellers who actually sell their tickets are small, the increase in prices is relatively modest and the probability of sale only falls slightly, so that overall platform revenues are significantly higher. Of course, this counterfactual is likely to be missing the fact that such large increases would likely cause some sellers to switch their sales efforts to other platforms, even if those

platforms have fewer buyers, in the last weeks before the game.

The final row in the table shows the proportional change in expected consumer surplus for a representative consumer arriving on each day (putting consumer welfare into dollar terms is complicated by the use of $\log(p)$ in the demand function) when the time-varying commission structure is adopted. Consumer surplus is affected by both prices and the availability of listings (e.g., low initial prices result in sales which reduces the availability of listings) but the changes show that the dominant effect is via price, with consumer surplus rising, by between 1% and 4% when prices fall, and falling when prices rise.

8 Conclusion

This paper sets out a framework for carrying out counterfactuals on platform design in environments where there are many competing sellers who face a price problem with a non-trivial dynamic component. This characterization applies to many online platforms which are trying to facilitate trade in perishable products, such as event tickets, accommodation and radio and television advertising. My framework allows for dynamics, competition and heterogeneity in both sellers and the products being sold, at the cost of making some strong simplifications about possible strategic behavior on the part of sellers and how market aggregates, here approximated by inclusive values, might evolve. In the empirical setting considered here, these simplifications seem appropriate given that most sellers are small and all sellers are anonymous, while prices tend to evolve in quite systematic ways at least until immediately before the event. Indeed one might also view the simplifications that are made as being reasonable approximations to the simplifications that sellers themselves are likely to use given the computational complexity that would be involved in solving for fully optimal strategies if all sellers were assumed to behave in an equally sophisticated way.

The counterfactual considered here is intended to primarily be illustrative of how the method works. In future iterations of the paper, I plan to investigate this type of counterfactual in more detail, and also to study the possible revenue and welfare effects of changing some features of demand. A couple of examples suggest themselves. One is recent changes on Stubhub (for example, as reported in the Wall Street Journal on September 1, 2015) that have moved the site between showing potential buyers prices that are inclusive of all fees and service charges and prices that do not, that are reported to have had both large effects on demand and large effects on platform revenues. An alternative example would be considering changes to the way that listings are presented. For example, demand on Stubhub

may be more price-elastic than demand on eBay (reported in Sweeting (2012)) partly because listings are, by default, shown with the cheapest ticket listed first, which may lead to consumers being more likely to purchase cheaper tickets if they use limited ‘consideration sets’ rather than considering the full set of tickets that are available. Using a different ranking system may therefore have the scope to increase platform revenues, seller revenues and buyer welfare simultaneously if it both makes demand less elastic but also matches potential buyers to listings more effectively. In both cases it would be interesting to both make predictions using the model, and to measure welfare effects, while also trying to validate the model by comparing its predictions to what actually happens when changes are made to the site. In particular this would make it possible to evaluate how important margins that are not included in the model (such as substitution of sellers between Stubhub and other sites) actually are.

References

- ARCIDIACONO, P., P. BAYER, J. R. BLEVINS, AND P. B. ELLICKSON (2012): “Estimation of dynamic discrete choice models in continuous time with an application to retail competition,” Discussion Paper 18449, National Bureau of Economic Research.
- 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.
- BENKARD, C. L., P. JEZIORSKI, AND G. Y. WEINTRAUB (2013): “Oblivious equilibrium for concentrated industries,” Discussion paper, National Bureau of Economic Research.
- COURTY, P. (2003): “Ticket Pricing under Demand Uncertainty*,” *Journal of Law and Economics*, 46(2), 627–652.
- DINERSTEIN, M., L. EINAV, J. LEVIN, AND N. SUNDARESAN (2014): “Consumer price search and platform design in internet commerce,” Discussion Paper 20415, National Bureau of Economic Research.
- DORASZELSKI, U., AND A. PAKES (2007): “A framework for applied dynamic analysis in IO,” vol. 3, pp. 1887–1966. Elsevier.
- ELLISON, S. F., AND C. M. SNYDER (2014): “An Empirical Study of Pricing Strategies in an Online Market with High-Frequency Price Information,” .

- GALLEGO, G., AND G. VAN RYZIN (1994): “Optimal dynamic pricing of inventories with stochastic demand over finite horizons,” *Management science*, pp. 999–1020.
- HAMMOND, R. (2008): “A Structural Model of Competing Sellers: Auctions and Posted Prices,” Discussion paper, mimeo, North Carolina State University.
- HÖRNER, J., AND L. SAMUELSON (2011): “Managing Strategic Buyers,” *The Journal of Political Economy*, 119(3), 379–425.
- LAZAREV, J. (2013): “The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in US monopoly markets,” Discussion paper, NYU.
- LEE, C.-Y., J. ROBERTS, AND A. SWEETING (2012): “Competition and Dynamic Pricing in a Perishable Goods Market,” Discussion paper, U Maryland.
- LESLIE, P., AND A. SORENSEN (2014): “Resale and rent-seeking: An application to ticket markets,” *The Review of Economic Studies*, 81(1), 266–300.
- MCASFEE, P., AND V. TE VELDE (2006): *Handbook on Economics and Information Systems*chap. Dynamic pricing in the airline industry. Elsevier.
- PAKES, A., AND P. MCGUIRE (1994): “Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model,” *The RAND Journal of Economics*, 25(4), pp. 555–589.
- QI, S. (2013): “The impact of advertising regulation on industry: The cigarette advertising ban of 1971,” *The RAND Journal of Economics*, 44(2), 215–248.
- SAEEDI, M. (2014): “Reputation and Adverse Selection, Theory and Evidence from eBay,” Discussion paper, Ohio State University.
- SWEETING, A. (2012): “Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets,” *Journal of Political Economy*, 120(6), 1133–1172.
- WEINTRAUB, G. Y., C. L. BENKARD, P. JEZIORSKI, AND B. VAN ROY (2008): “Nonstationary Oblivious Equilibrium,” Discussion paper, Columbia University.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): “Markov perfect industry dynamics with many firms,” *Econometrica*, 76(6), 1375–1411.

XU, D. Y. (2008): “A structural empirical model of R&D, firm heterogeneity, and industry evolution,”
Discussion paper, New York Univ.