Staggered vs. Simultaneous Price Setting with an Application to an Online Market

Andrew Sweeting* Kane Sweeney

University of Maryland eBay Research Labs

March 9, 2015

Abstract

The typical assumption in structural I.O. is that sellers set prices or quantities simultaneously. However, many markets, including e-commerce markets, may be better approximated by a continuous-time type of framework where a seller can change its prices instantaneously but does so assuming that its rivals are not changing their prices at the same moment. We illustrate how this can affect the level and dynamics of equilibrium prices, even when there are a number of competitors. We apply our model to the empirical setting of secondary market pricing of event tickets on Stubhub.com to understand whether the nature of price-setting plays a role in determining the dynamics of prices as an event approaches.

*Contact: sweeting@econ.umd.edu. This draft is preliminary and incomplete. This is paper developed out of work sponsored by NSF Grant SES-0850441 (Sweeting). Siyuan Tu, Chung-Ying Lee and Xuezhen Tao provided excellent research assistance. Stubhub provided access to their data, and we are grateful to an anonymous ticket broker for useful discussions.

1 Introduction

Understanding the details of how prices are set is one of the most important contributions that I.O can make to economics (Carlton (1989)). However, most applied work makes a very strong simplifying assumption that sellers set prices simultaneously, and usually statically, even when allowing for each dynamics in capacity and investment choices. While simultaneous price setting is a convenient assumption, and undoubtedly appropriate in some settings, such as sealed-bid auctions or health insurance markets where firms design contracts that will be offered throughout an open enrollment period, it is clearly a simplification even in markets where firms can adjust their prices instantaneously, but because of limited managerial attention or because of idiosyncratic and time-varying costs of adjustment, it is very unlikely that two firms will change their prices at the same moment. As has been shown in macroeconomics, usually in the context of models with limited strategic interactions[1], non-simultaneous price-setting can have important impacts on business cycle dynamics and adjustment to both domestic and international shocks (Taylor (1980), Calvo (1983) and Caplin and Spulber (1987); and recent surveys by Klenow and Malin (2010) and Nakamura and Steinsson (2013)), and it is clearly potentially important to investigate how more realistic models of price setting could impact outcomes of more traditional interest to I.O. economists, taking into account the types of strategic interactions that can matter in these settings. In particular, a seller might be willing set a price that is significantly above static Nash equilibrium levels if this will tend to raise the prices that its competitors set in future periods.

In this paper, we try to understand the role of non-simultaneous price-setting in explaining the dynamics of prices in a secondary market for event tickets on Stubhub.com. As documented in Sweeting (2012), the typical pattern in these markets is for prices to decline quite dramatically and predictably (e.g., 30-40% in the final month) as the event approaches. In that paper, each seller is assumed to be so small that he ignores any effect his price may have on the prices set by other sellers. Under this assumption, the price declines are explained as coming from the perishability of the products, as sellers become increasingly keen to sell as the date when their tickets lose their value approaches. However, while these assumptions are reasonable for the 2007 eBay market for Major League baseball tickets which was Sweeting's focus[2], there may not hold in other markets

[1]In the sense that a firm will not consider how the price that *it* sets will *change* the future prices set by other sellers, even if it is fully rational in knowing that those prices will adjust over time.

[2]MLB teams play 81 regular season home games, and season ticket holders who do not want to go to a particular game account for a large number of tickets. Therefore a model where there are a lot of small sellers of one unit is appropriate. Second, the focus of Sweeting's work was the eBay market in 2007 where the ability of both buyers and

even when very similar products are being sold. Indeed, our interest in this topic was partly motivated by discussions with a large ticket broker - whose pricing behavior we shall look at below - who accounts for a large proportion of the tickets listed and sold in the Stubhub market (which in general is several times larger than eBay) for some events and who was well aware of the fact that when it cuts prices, it tends to lower the prices set afterwards by its smaller competitors. The broker was particularly concerned about this ability to 'move the market' because demand in these markets is greatest closer to the event, so that an early price cut may generate only small increases in revenues at the time that the price falls, while proving costly if it leads to lower market prices when more consumers are purchasing. In this paper, we try to quantify the importance of non-simultaneous price setting in the dynamics of prices, when compared to perishability and possible changes in the elasticity of demand as an event approaches, focusing on events where this broker was a particularly large seller. Of course, we believe that there may be many other settings where non-simultaneous price setting is important and we discuss some of these contexts in the conclusion.

While applied work has largely assumed simultaneous price-setting, even in the context of rich, dynamic models (e.g., Ericson and Pakes (1995), Pakes and McGuire (1994) and applications such as Ryan (2012),Snider (2009) and Jeziorski (2014)), Maskin and Tirole's initial analyses of games with Markov Perfect Nash Equilibria (MPNE) (Maskin and Tirole (1988a), Maskin and Tirole (1988b) (MTII hereafter) and Maskin and Tirole (1987)), focused on games where homogeneous goods duopolists set prices or quantities in alternating periods in infinite horizon games with time-invariant demand and marginal costs. When a rival firm is making its price choice, a firm is committed to the price or quantity it chose in the previous period. They show that in any MPNE of this staggered price setting game, firm profits will be above the (zero) level that they would have if firms set prices simultaneously and they played the static Nash equilibrium, and that two types of MPNE exist. In a focal price equilibrium, both firms set the same price forever and this price is substantially above the competitive price.[3] The unique renegotiation-proof MPNE, when firms are sufficiently patient, is a focal price equilibrium at the (static) monopoly price. Equilibrium prices above marginal cost are supported by the fact that, while a firm could increase its short-

sellers to search and compare ticket prices was fairly poor, so that strategic interactions between similar tickets were likely to be weak. In an empirical analysis, Sweeting could not reject no strategic effect. In contrast, comparing similar tickets (e.g., those in the same section) on Stubhub is straightforward and there is evidence of significant interactions.

[3] In MTII's model, a focal price equilibrium must involve the firms securing at least $\frac{4}{7}$ths of the monopoly profit.

run profits by undercutting its rival, it believes that this will result in itself being subsequently undercut, driving future prices down to a lower level. While pricing at the monopoly price could also be supported as the outcome of tacit collusion in the supergame equilibrium of a simultaneous move game when firms are patient enough, the incentives to charge monopoly prices in the MPNE are rather different and instead rely on the fact that firms dynamic pricing strategies are, at least for some subsets of prices, strategic complements.[4] The other type of equilibrium they identify is one, which they call an Edgeworth cycle, in which prices move in a continuous cycle, with firms undercutting each other until prices become so low that one of the firm chooses to a high price, so that undercutting can start again but from a higher price level.

Applied research based on MTII has primarily focused on Edgeworth cycles, in particular trying to understand how closely the price cycles that are observed in many retail gasoline markets (Eckert (2013), Eckert and West (2004), Noel (2007), Noel (2008), De Roos and Katayama (2013), Wang (2009)) conform to the predictions of MTII's theory, allowing for the facts that, unlike in MTII, gasoline stations are spatially differentiated, possibly capacity constrained, face a variety of cost and demand shocks and, in some jurisdictions, are actually forced to set prices simultaneously because of regulation. However, relatively little attention has been given to the fact that in many other markets, with the exceptions of those which are cleared through centralized auction mechanisms, firms do not set prices at the same time and so MTII's theory might suggest that prices might diverge significantly from competitive levels even if we do not observe regular price cycles and even if firms do not use collusive strategies. However, because existing work has focused entirely on single-product duopoly or occasionally triopoly settings (e.g., Noel (2008) who computes equilibria in MTII models extended in various directions, and shows that cycles are a fairly robust phenomenon, while maintaining the infinite horizon assumption), it is unclear how large the effects of non-simultaneous price setting can be in markets that are more competitive. In this paper we extend this literature by considering both examples and an empirical setting with more sellers. We show that prices

[4]A couple of comparisons help to illustrate the differences. First, in a supergame where firms are sufficiently patient monopoly outcomes can be supported if firms choose quantities rather than prices. In contrast, in the sequential move game, quantity-competition with patient firms results in outcomes that are more competitive than static Cournot, as an increased quantity by one firm will tend to reduce the amount that its rival wants to produce when it gets to adjust, therefore increasing competitiveness. Second, as we illustrate below, above static equilibrium prices can be supported, at least for a substantial number of periods, in a finite horizon version of the sequential move game whereas only static Nash strategies form a sub-game perfect equilibrium when firms move simultaneously (see also Wallner (1999) who shows this result for the duopoly case). Note that while in a static, simultaneous move model prices will be strategic complements, in the alternating move model a firm may not want to increase its prices when its rival price is low enough. Therefore in some regions prices may be dynamic strategic substitutes rather than complements. We illustrate the shapes that reaction functions can take below.

can be significantly above static, simultaneous Nash levels even as the number of sellers grows.[5] Contrary to the existing literature we focus on finite horizon games. This has the advantage that we are able to prove the uniqueness of our equilibrium via backwards induction.

We will discuss related work on the dynamic pricing of perishable goods below when we lay out a theoretical framework for our analysis. The paper is also related to the recent literature on dynamic games set in continuous time (Doraszelski and Judd (2012),Arcidiacono, Bayer, Blevins, and Ellickson (2012)), although applications have been limited (an exception is Jeziorski (2014), who models radio station mergers and format choices in continuous time, showing that the framework can handle ambitious problems that could not be addressed in earlier models such as Sweeting (2013)). This literature, which considers stationary, infinite horizon games, emphasizes the computational advantages of continuous time that result from the fact that, at least when all moves are observed instantaneously, the probability that the state of more than one player changes can be ignored. This simplifies the integration over future states required when calculating value functions, and this is especially important when choice sets are large (as they must be when considering realistic price choices). While these ideas partly inspire the approaches we take in this paper, we gain a similar advantage by using discrete time intervals that we assume are 'short enough' that the probability that two players would move at the same time is small enough to be ignored, as an approximation of a continuous time game. The reason why we do not formally consider continuous time is that the convenient formulae that exist for calculating seller values in stationary, infinite horizon games, cannot be used in our finite horizon setting.

As noted above, the macroeconomics literature has explored non-simultaneous price changes. One approach assumes that staggered price-setting as a primitive assumption of the model (Calvo (1983)), whereas in other models it emerges in equilibrium from firms having idiosyncratic menu costs even though firms always have the ability to change their prices (Caplin and Spulber (1987)). In this paper we assume that a price change opportunity arrives with some probability for each seller in each of our short periods (and that nature will select only one player), similar in spirit to Arcidiacono, Bayer, Blevins, and Ellickson (2012), but that it is costless to change prices once this opportunity arises. Our approach is similar to the one that Ellison and Snyder (2014) suggest as

[5] This being said, computational limits impose restrictions on the number of firms that can be considered. While it is tempting to use approximate dynamic programming techniques, such as trying to approximate value functions by projecting onto a set of basis functions (e.g., Rust (2000), Arcidiacono, Bayer, Bugni, and James (2012)) to extend the results to even larger markets, the fact that firms' reaction functions can be discontinuous means that these types of approximations may work very poorly.

being appropriate for the online market for memory chips, where they observe a set of small firms changing their prices that appear on the price search engine 'Pricewatch.com'. They argue that an appropriate model to explain pricing behavior is one where managers are inattentive to what is changing on Pricewatch and then, occasionally, check to see if their current price is appropriate and, if not, update it appropriately.[6] They examine the strategies of different types of firm but they do not, unlike us, attempt to solve for what should happen in equilibrium. The few papers that have attempted to solve for price equilibria models of competing sellers in online markets have typically assumed simultaneous price-setting behavior. Recent examples include Dinerstein, Einav, Levin, and Sundaresan (2014) and Zhu (2014) (who also considers a Stubhub market for event tickets).

The paper proceeds as follows. Section 2 lays out a general theoretical framework for the dynamic pricing of perishable goods, and Section 3 provides example of the type of dynamics that can be generated by non-simultaneous price setting without perishability. Section 4 describes the empirical setting that we consider, discussing whether existing models of oligopolistic perishable goods pricing with simultaneous price setting explain the observed pricing behavior. We also argue, by focusing on the large broker mentioned before, that a seller can move the market when its sets its price, suggesting that sellers should be thinking about their price setting strategically. Section 5 presents our full model and details the estimation of several key parameters. Sections 6 and 7 (to be completed) will present our counterfactuals and conclusions. All results in the current draft should be regarded as preliminary.

2 Theoretical Framework for the Pricing of Perishable Goods

This section presents a simple theoretical framework to help understand sellers' pricing decisions and to explain the relationship between the analysis in the current paper and the existing revenue management literature on the dynamic pricing of perishable goods, whether under monopoly or some form of competition.

Suppose that a seller i at time t has n_{it} units of a perishable product to sell (the expiration date and the end of the game is T), and wants to set an optimal price p^*_{it}. Assume that time periods are short so that if it sets price p_{it} the probability that it sells one unit is $q_{it}(p_{it}, \mathbf{p}_{-it}, \mathbf{H}_t)$ where

[6] They also discuss how similar what the pricing patterns that they observe are to Edgeworth cycles. In a classic MTII-style Edgeworth cycle, a one firm occasionally *raises* its prices, and this is followed by a sequence of price changes where firms undercut each other. In contrast, what they observe is that firms occasionally *drop* their prices, so that they are ranked much more favorably on Pricewatch, but that this is then followed by other firms undercutting them.

$\mathbf{H}_t$ collects state variables that can also affect demand (e.g., the number of other seller listings and their characteristics, the characteristics of buyers who may show up) and $\mathbf{p}_{-it}$ are the prices set by other sellers, with zero probability that two units or more units are sold at time t. We assume that $\frac{\partial q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\partial p^*_{it}} < 0$, for $\forall i, t$. i's (time t) expected value to holding n units of inventory at time $t+1$ is $\mathbb{E}_t(V_{it+1}(n, \mathbf{H}_{t+1})|p_{it}, p_{-it}, \mathbf{H}_t)$. i will choose a strategy to maximize its value, defined by the Bellman equation:

$$V_{it}(n_{it}, \mathbf{H}_t) = \max_{p_{it}} p_{it} q_{it}(p_{it}, p_{-it}, \mathbf{H}_t) + q_{it}(p_{it}, p_{-it}, \mathbf{H}_t)\mathbb{E}_t(V_{it+1}(n_{it}-1, \mathbf{H}_{t+1})|p_{it}, p_{-it}, \mathbf{H}_t) + ... \quad (1)$$
$$(1 - q_{it}(p_{it}, p_{-it}, \mathbf{H}_t))\mathbb{E}_t(V_{it+1}(n_{it}, \mathbf{H}_{t+1})|p_{it}, p_{-it}, \mathbf{H}_t)$$

and assuming that the demand is downward sloping and that both the demand and value functions are differentiable, the optimal price will be:

$$p^*_{it} = \{\mathbb{E}_t(V_{it+1}(n_{it}, \mathbf{H}_{t+1})|p^*_{it}, p_{-it}, \mathbf{H}_t) - \mathbb{E}_t(V_{it+1}(n_{it}-1, \mathbf{H}_{t+1})|p^*_{it}, p_{-it}, \mathbf{H}_t)\} + ...$$
$$\frac{q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\left|\frac{\partial q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\partial p^*_{it}}\right|} + \frac{q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)\frac{d\mathbb{E}_t(V_{it+1}(n_{it}-1, \mathbf{H}_{t+1})|p^*_{it}, p_{-it}, \mathbf{H}_t)}{dp_{it}}}{\left|\frac{\partial q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\partial p^*_{it}}\right|} + ... \quad (2)$$
$$\frac{(1 - q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t))\frac{d\mathbb{E}_t(V_{it+1}(n_{it}, \mathbf{H}_{t+1})|p^*_{it}, p_{-it}, \mathbf{H}_t)}{dp_{it}}}{\left|\frac{\partial q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\partial p^*_{it}}\right|}$$

The first term, the difference between the seller's expected value in the next period when it has only $n_{it} - 1$ units rather than n_{it} units, represents the (marginal) cost of a sale.[7] The second term $\left(\frac{q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\left|\frac{\partial q_{it}(p^*_{it}, p_{-it}, \mathbf{H}_t)}{\partial p^*_{it}}\right|}\right)$, is the standard (static) mark-up that depends on the level and slope of current demand. The final two terms reflect how a seller's current price may affect its future values (other than the direct effect through the sale of the unit). We have shown these as total derivatives to reflect the fact that p_{it} may increase values through a number of different mechanisms.

Before discussing them, consider the existing dynamic pricing literature. In the classic dynamic pricing papers of Gallego and Van Ryzin (1994) and McAfee and Te Velde (2006), there is a single (monopoly) seller; potential buyers must purchase at once or exit the market forever, so that there is no scope for them to respond to high current prices by strategically delaying purchase, increasing

[7] In a monopoly problem with free disposal and non-strategic buyers the seller's value should be increasing in the number of units it has to sell. However, this is not necessarily true when there are multiple sellers or buyers are strategic.

future demand; and, the seller can change prices every period. Therefore, the $\frac{d\mathbb{E}_t(V_{it+1}(.))}{dp_{it}}$ terms are zero. In this case, the expected path of prices is driven by the intertemporal trajectory of $\{\mathbb{E}_t(V_{it+1}(n_{it}, \mathbf{H}_{t+1})|p^*_{it}, p_{-it}, \mathbf{H}_t) - \mathbb{E}_t(V_{it+1}(n_{it}-1, \mathbf{H}_{t+1})|p^*_{it}, p_{-it}, \mathbf{H}_t)\}$ (sometimes, called the 'bid-price') and changes in demand. Pang, Berman, and Hu (2015) show that, under optimal pricing, the expected path of the bid-price is upwards until only one unit remains for sale, and that therefore, without changes in demand, the expected price of a seller with multiple units tends to increase until close to the end of the horizon. This is consistent with the simulations in Gallego and Van Ryzin (1994) and McAfee and Te Velde (2006) where expected prices tend to be slightly increasing until close to T. As initial inventories become very large the expected price is close to flat, and the gains to dynamic pricing (without demand changes) become minimal. On the other hand, the prices of sellers with a single unit should tend to monotonically fall. Of course, changes in a seller's residual demand will also tend to affect mark-ups.[8]

In more general models, the $\frac{d\mathbb{E}_t(V_{it+1}(.))}{dp_{it}}$ terms may not be zero. There are a number of possible ways in which a seller's current price can affect its future value. First, buyers may be strategic, so that they delay purchasing when prices are expected to fall or variety to increase[9] As demand is pushed into the future, $\frac{d\mathbb{E}_t(V_{it+1}(.))}{dp_{it}} \geq 0$, which will tend to increase the current optimal price, but it will also tend to make current demand more elastic, depressing the optimal price. Considering a monopolist and assuming no ability to commit to a future price path, Levin, McGill, and Nediak (2009) show that the introduction of strategic consumers tends to flatten the path of prices relative to the case where consumers are not strategic.[10]

Second, in a setting with multiple sellers, i has an incentive to set a high current price so that other sellers are more likely to sell their inventory, which will raise competitor prices closer to T. However, because these future competitor prices will also depend on i's future inventory, the situation is actually more ambiguous. For example, Gallego and Hu (2007) characterize the open-loop

[8]One could prove some potentially testable results in this model. For example, suppose that two sellers will face the same residual demand in every period but which may change over time, but that one seller has multiple units to sell and the other has one; and that both of them have no value to units after T. Then, the expected price path - conditional on not selling - should be flatter for the multi-unit seller (specifically, optimal prices would be the same in the final time period, but lower in previous periods). This prediction would be rejected by our data, but this could be explained by the one unit seller having a much higher value of being left with a ticket at the time of the event than a large ticket broker.

[9]In our setting, the number of available listings and most measures of variety fall as an event approaches.

[10]Horner and Samuelson (2011) consider this question in a setting where the seller has only one unit to sale and where all potential buyers are in the market from the beginning. They show that the seller will either maintain very high prices, and make no sales, until right at the end when it will drop the price discretely, or alternatively it will drop the price smoothly, as in a Dutch auction. The optimal policy depends on the number of buyers.

Nash equilibrium in an differential oligopoly game with differentiated products and deterministic demand. They show prices will tend to increase over time on the equilibrium path as sellers initially price aggressively in order to reduce their inventories.[11] Lin and Sibdari (2009) consider a discrete time duopoly game with logit demand, where, for a given level of inventory, prices show a non-monotonic pattern with the time remaining, as, at intermediate time horizons, the incentive of the firm to try to reduce its own inventory to raise future prices dominates. All of these models assume that firms either adjust prices continuously (e.g., Gallego and Hu (2007)) or at least simultaneously.

Third, when firms do not set prices at the same time (for example, sellers only occasionally consider when to adjust their prices), an additional incentive to set a higher initial price can be created if doing so raises the price set by sellers who set their prices subsequently, even when these sellers recognize that this may allow them to be subsequently undercut. In this setting, the current (standing) prices of other firms, set in previous periods, will enter the policy function of a seller that is currently changing its price. When prices are dynamic strategic complements (so that a seller's optimal price is increasing in the prices previously set by other sellers), this provides an incentive for these earlier sellers to raise their prices. Therefore this potentially reinforces the incentives of early price setters to set high initial prices. Of course, as the end of the time horizon approaches, the incentive to try to raise future prices disappears. This is true, with or without perishability, as the following computational examples illustrate.

3 Example with Non-Simultaneous Price Setting in a Finite Horizon Game

Consider a finite-horizon, discrete time game with T periods (in our examples $T \geq 2,000$), where we consider each period to be short enough so that only one 'event' can happen. In the base case, there is no discounting (motivated by the fact that we are considering relatively short time horizons). There are N symmetric firms, each of which chooses prices from a discrete grid $[\underline{p}, \overline{p}]$. The firms sell a good that is costless to produce when someone wants to purchase it, so that there is no perishable inventory and a seller remains in the market when a sale is made. The point of the examples are to show what happens to prices when price setting is not simultaneous, without the dynamics that are created by perishability. In each period, two mutually exclusive (and non-exhaustive) types of

[11] Perakis and Sood (2006) also predict increasing prices in a oligopoly model where firms set both prices and 'protection levels' (i.e., limits on the amount of inventory that can be sold at given points in time).

events can happen.

1. a buyer arrives in the market with probability $q^{\text{buyer}}(t)$. The buyer buys one unit of the good for sale from at most one seller. The preferences of the buyer are described by a nested logit model where all of the N sellers are grouped into a single active nest, and the buyer's utility from purchasing from seller j is:

$$u_j = v - \alpha p_j + (1 - \sigma)\varepsilon_j, \tag{3}$$

   and the utility of not purchasing is equal to 0. In the examples below we set $v = 3, \alpha = 1$ and the nesting parameter can take values $\sigma = 0.5$ or 0.9. In the latter case, the products are close substitutes and, a seller's market share will increase significantly when it has the lowest price. Note that buyers behave statically in the sense that they only care about the current price. A seller remains in the market when a sale is made, so that, even though the time horizon is finite, the good is not perishable.

2. a seller can change its price. We contrast two situations. In one of them, all sellers get to change their price simultaneously, which happens with probability $q^{\text{p change sim}}(t)$. In the other, only one seller gets to change its price, with this seller chosen with probability $q^{\text{p change stag}}(t)$. The probabilities are independent of the prevailing prices, and, when the opportunity to move arises, changing price from its current level is costless.

In each time period, nature moves first, by choosing what event will happen (including which seller will gets to change its price), followed by the purchase decision or price change decision being made. If no event occurs (which is the most likely outcome given the chosen parameters) the game moves onto the next period. Note that in this specification there are no payoff shocks that affect the level of the optimal price that a player sets in any period.[12]

3.1 Solution

Our focus is on the expected path of prices in this model. We solve for the unique subgame perfect equilibrium pricing policies using backwards induction.[13] A policy consists of the price that a firm

[12] In particular, it would be interesting to test whether the introduction of shocks (which one might view as temporary shocks to marginal costs) would (i) substantially affect the relationship between the effects of staggered pricing and the number of players, and (ii) tend to smooth out the path of the (average) optimal price a long way from the end of the time horizon, which in the currently model can exhibit quite an unstable pattern.

[13] In this setting the Markov Perfect Nash equilibrium will be identical to the subgame perfect Nash equilibrium.

will set if nature allows it to change its price. When firms set prices simultaneously, the usual logic implies that the unique equilibrium will involve firms setting the (unique) static Nash equilibrium prices in any period when they change price.[14] When price setting is staggered, a seller's pricing policy will be conditional on other seller's current prices (which are, of course, payoff relevant because they will affect the demand of the price-changing seller if a buyer shows up in subsequent periods). A player's value in period t is defined as its value at the beginning of the period, when its price is p_{it} and the prices of other sellers are $\mathbf{p}_{-it}$ (**bold** denotes vectors). A player's optimal policy is then

$$p^*_{it} = \max_{p'_{it}} V_{t+1}(p'_{it}, \mathbf{p}_{-it}) \tag{4}$$

and its period t value, which is defined at the beginning of the period before nature gets to choose what event will occur, can be calculated as

$$V_t(p_{it}, \boldsymbol{p}_{-it}) \equiv q^{\text{buyer}}(t)(p_{it}) \Pr_i\left(\text{sale}_i | p_{it}, \boldsymbol{p}_{-it}\right) + q^{\text{p change sim}}(t) V_{t+1}(p^*_{it}, \boldsymbol{p}_{-it}) + \ldots \tag{5}$$
$$\sum_{j \neq i} q^{\text{p change sim}}(t) V_{t+1}(p_{it}, p^*_{jt}, \boldsymbol{p}_{-\{i,j\}t}) + \left(1 - q^{\text{buyer}}(t) - \sum_{j \neq i} q^{\text{p change sim}}(t)\right) V_{t+1}(p_{it}, \boldsymbol{p}_{-it})$$

Solving the game as N increases is computationally demanding, with the computation time a function of the richness of the price grid. For example, when $N = 7$ and we consider a price grid with 24 possible prices, there are 11,400,480 states even when we impose symmetry and exploit exchangeability. The computation time associated with this number of states is significant because, even when only one player can move, it is possible to transition from any state to a large number of possible states in the following period.[15,16] Efficient solution therefore requires pre-computing

[14] Note that this is true because we include the static Nash equilibrium price from a game with continuous price choices as one of the points on our price grid. In our reported solutions we use a grid with a small number of points beneath the static Nash equilibrium price, although the expected price paths are the same - and do not involve below static Nash prices - when we include more points below the static Nash price.

[15] Note that in many investment and capacity games involving large state spaces it is often assumed that a player's state variable (e.g., product quality or capacity) can increase or decrease by at most one unit in one period, which substantially limits the number of possible transitions that can occur. In a price setting game, this type of restriction does not make sense, so the number of possible transitions remains large. Also in the finite horizon game it is necessary to compute value functions and policies in each period of the game, whereas in an infinite horizon, but stationary, game it is only necessary to compute a single equilibrium value and policy function.

[16] The richness of the price grid can matter for the results, as well as the plausibility of the model (given that online prices are fairly closer to being continuous). Through computing a large number of examples, considering either a finer grid or a grid with a higher upper bound, or both, tends to increase how far prices lie above the static Nash equilibrium level early in the time horizon. This suggests that the price levels reported below may underestimate the prices that can be sustained with sequential price setting. This issue will be studied more thoroughly in later iterations of the paper.

indices for the possible transitions, so that values can be easily compared when computing policies.

While it is possible to calculate a unique set of equilibrium strategies using backwards induction, to simulate price paths one needs to assume which state the market is in the initial period. Here we consider a couple of alternatives and show that, after the early part of the game, the chosen alternative has very little effect on the average level of prices. The alternatives are:

1. prices start at the (symmetric) static Nash equilibrium price level; or,

2. prices start at the equilibrium level which firms would choose given their value functions in the first period of the game that we consider (i.e., what would be the Nash equilibrium in a fictitious one-shot game before the actual game begins where there is no demand and firms would choose their initial prices for the game.[17]

At least for the parameters that we consider, these alternative starting points lead to the same price paths after some initial set of periods as prices adjust. In the results that follow we use the results from the second starting point.

3.2 Results

Table 1 presents the results from a set of different models where we assume that sellers are symmetric, the arrival rate of buyers into the market is fixed over time - the probability that a buyer arrives is equal to 0.02 each period - but we allow the number of sellers, the frequency with which they can change prices and the intensity with which they compete, determined by the nested logit parameter, σ, to vary. The left-hand side of the table shows results for $\sigma = 0.5$ and the right-hand side shows results for $\sigma = 0.9$. For each game, we solve for policies in each period and simulate 500 times for T =2,000 periods. The table reports: (i) the price that the sellers set in every period when they set prices simultaneously (this is just the static Nash equilibrium price); and when sellers set prices at different times (ii) the average (across simulations) price in the final time period, $t = T$; the average price 100 periods before the end ($t = T - 100$); (iii) the maximum price across the 2,000 periods; and, (iv) the average price across the 2,000 periods. The numbers in parentheses are the standard deviations in the reported statistic across the 500 simulations.

[17] We use the value functions in the first period of the game to first look for the symmetric pure strategy Nash equilibrium with the lowest prices, and if we do not find a symmetric PSNE we look for an asymmetric one, and, if there are several, we use the one with the lowest average price. If there is no PSNE of any sort (which we have not noticed happening in any of our models), we choose a symmetric starting point that minimizes the payoff incentive that any individual player has to deviating).

From the table, it is immediately apparent that prices on the equilibrium path, in at least some periods, with non-simultaneous price setting can be substantially above those that result from simultaneous price setting. As a comparison, the monopoly price for either value of the nesting parameter would be 2.56. The differences in prices are surprisingly large even when there are as many as seven players. For example, with $\sigma = 0.5$ and $N = 7$, average prices are similar or above the static Nash equilibrium level with 3 players, and maximum prices are significantly above the static Nash equilibrium price level with just two players. While all prices are much lower when $\sigma = 0.9$, fairly similar types of comparisons can be made. The additional mark-ups tend to disappear as the end of the time horizon approaches, so that by T prices are close to static Nash equilibrium levels. This reflects the fact that, as the end of the game approaches, the incentive to undercut rivals to gain sales immediately dominates an incentive to maintain higher prices so that rivals do not undercut. This is an interesting finding given that one of the patterns that we will want to explain is why prices on Stubhub decline so much as a game approaches.

To understand this logic, it is useful to look at the price path and reaction functions in Figure 1. In this case $N = 2$, $\sigma = 0.5$ and the probability that a seller gets to change its price is 0.01 in each period (so the results correspond to the first row of the left-hand side of the table). The top panel shows the path of a seller's expected price during the game (note the x-axis measures the number of periods from the end of the game, so the final period is on the left of the graph). The lower panels show a seller's reaction function (if it gets to change its price) in four different time periods, with the red line being the 45^o line so that any point below the red line is associated with a seller undercutting its rival. At $t = T - 10$ the reaction function resembles a seller's reaction function in a one-shot static game, as a seller's best response is to undercut at all prices above static Nash prices. As one moves further away from the end of the game the reaction functions change, with a seller being more willing to set a higher price in response to a particular price by its competitor, because this will influence the price that its rival sets if it is the next one to change its price. At $t = T - 100$, where the expected price is falling, the incentive remains to undercut, but by a smaller amount than at $t = T - 10$. The reaction functions at $t = T - 400$ (when expected prices are rising) and $t = T - 1000$ (when they are flat) have quite different shapes. In particular, a seller may respond to a low price from a competitor by setting a higher price, which can result, once the competitor makes a best response, by both firms setting prices that are substantially above the static Nash equilibrium level. When $t = T - 400$ the reaction function has a particularly non-monotonic form, with the particular peaks and troughs being able to support rising prices, substantially above Nash

Figure 1: Price Paths and Reaction Functions for $N = 2$, $\sigma = 0.5$, Price Change Probability 0.01

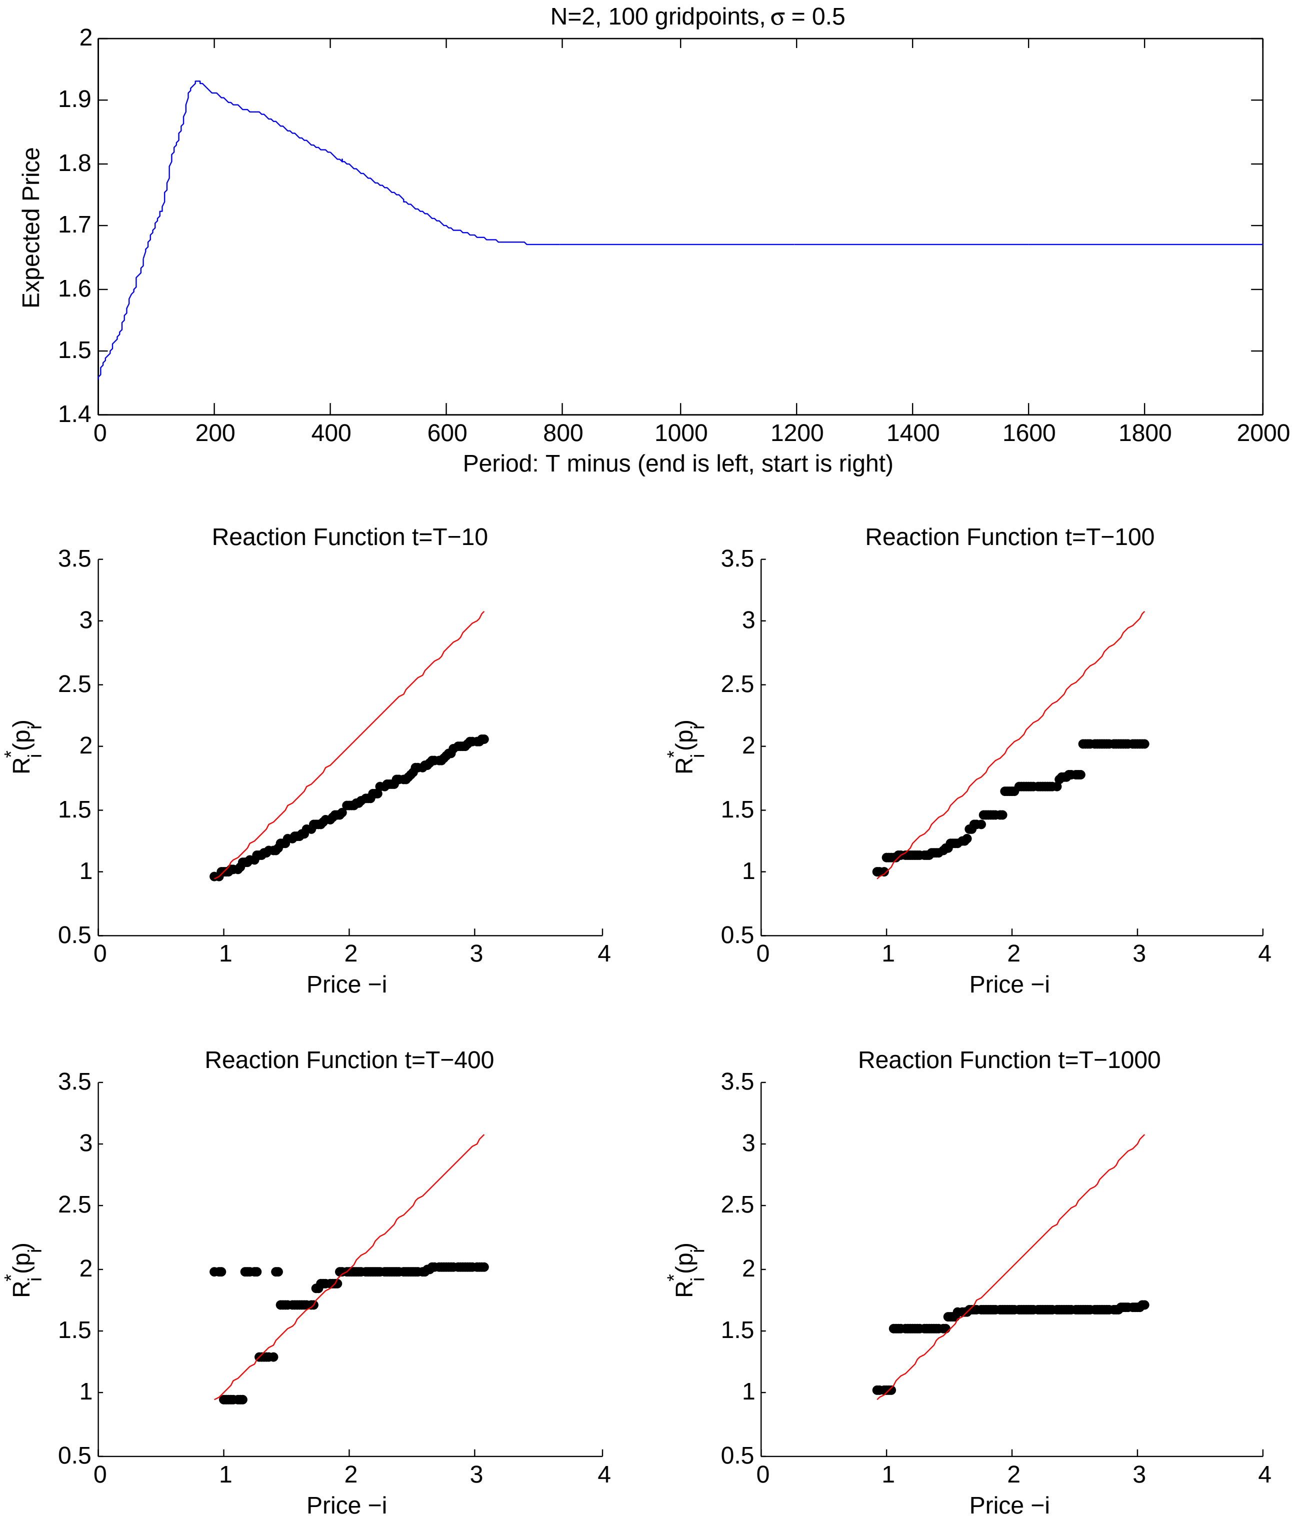

levels, as the equilibrium outcome.[18] Intuitively, the incentive to try to raise prices at this stage of the game comes from sellers knowing that, in the last part of the game, which at this point is about to begin, the dominant incentive will be to undercut. In this case, it is attractive to try to increase the market price in future periods even if a seller is sacrificing some sales when it raises its own current price. Note however, that this kind of strategy emerges without the type of supergame incentives that would exist in an infinite horizon setting.

As the number of players increases, the time at which prices tend to drop move further back from the end of the game, which is one reason why average prices tend to be lower, relative to the maximum prices, with more players. We note, however, that it is not necessarily the case that prices start high and then drop monotonically. Instead, depending on the parameters, and also the richness of the price grid considered, there can be sets of periods quite some time before the end of the game where prices tend to cycle upwards and downwards, although, unlike in the infinite horizon duopoly games considered by MTII, these cycles do not display regular features that keep on repeating.[19]

An example of this type of price path can be seen from the black line in Figure 2(a), drawn for the case where $N = 3$ and the probability that a player gets to change its price is 0.01 each period for each player, with $\sigma = 0.5$ (i.e., this corresponds to the fourth row on the left-hand side of Table 1). In this case, average prices rise steadily, peak around period $t = T - 800$, and then fall slightly before rising and then starting a steady descent to the end of the game.

As noted previously, one feature of our empirical setting is that demand is much higher towards the end of the game. While raising the probability that a buyer arrives equally across all periods does not affect equilibrium prices (the revenues from each price combination are simply multiplied by a constant, and in our model, where we do not have any costs of changing prices, the optimal price choice is therefore unchanged), raising the probability that a buyer arrives in some periods while reducing it in others does affect the path of prices. As an illustration, Figure 2(a) also shows the path of prices when we change the pattern of demand in two ways. The red line is associated with higher demand in the last 400 periods of the game, while the blue line is associated with

[18] Partly for this reason, patterns with rising prices tend to be more sensitive to the definition of the price grid. Based on a number of experiments conducted so far, stretches where prices steadily rise appear to be easier to support as the price grid becomes richer and longer. For this reason, one might expect this to arise in more realistic models where prices are close to continuous.

[19] Note that the price paths are also different from what emerges in a finite horizon game with alternating, but certain, price setting opportunities. In the duopoly case Wallner (1999) shows that the only possible equilibrium will involve cycles, and will not involve periods of flat prices. In contrast, as can be seen in the previous figure, we do find periods of constant prices in our model.

Figure 2: Price Paths for $N = 3$, $\sigma = 0.5$, Price Change Probability 0.01, Time-Varying Demand

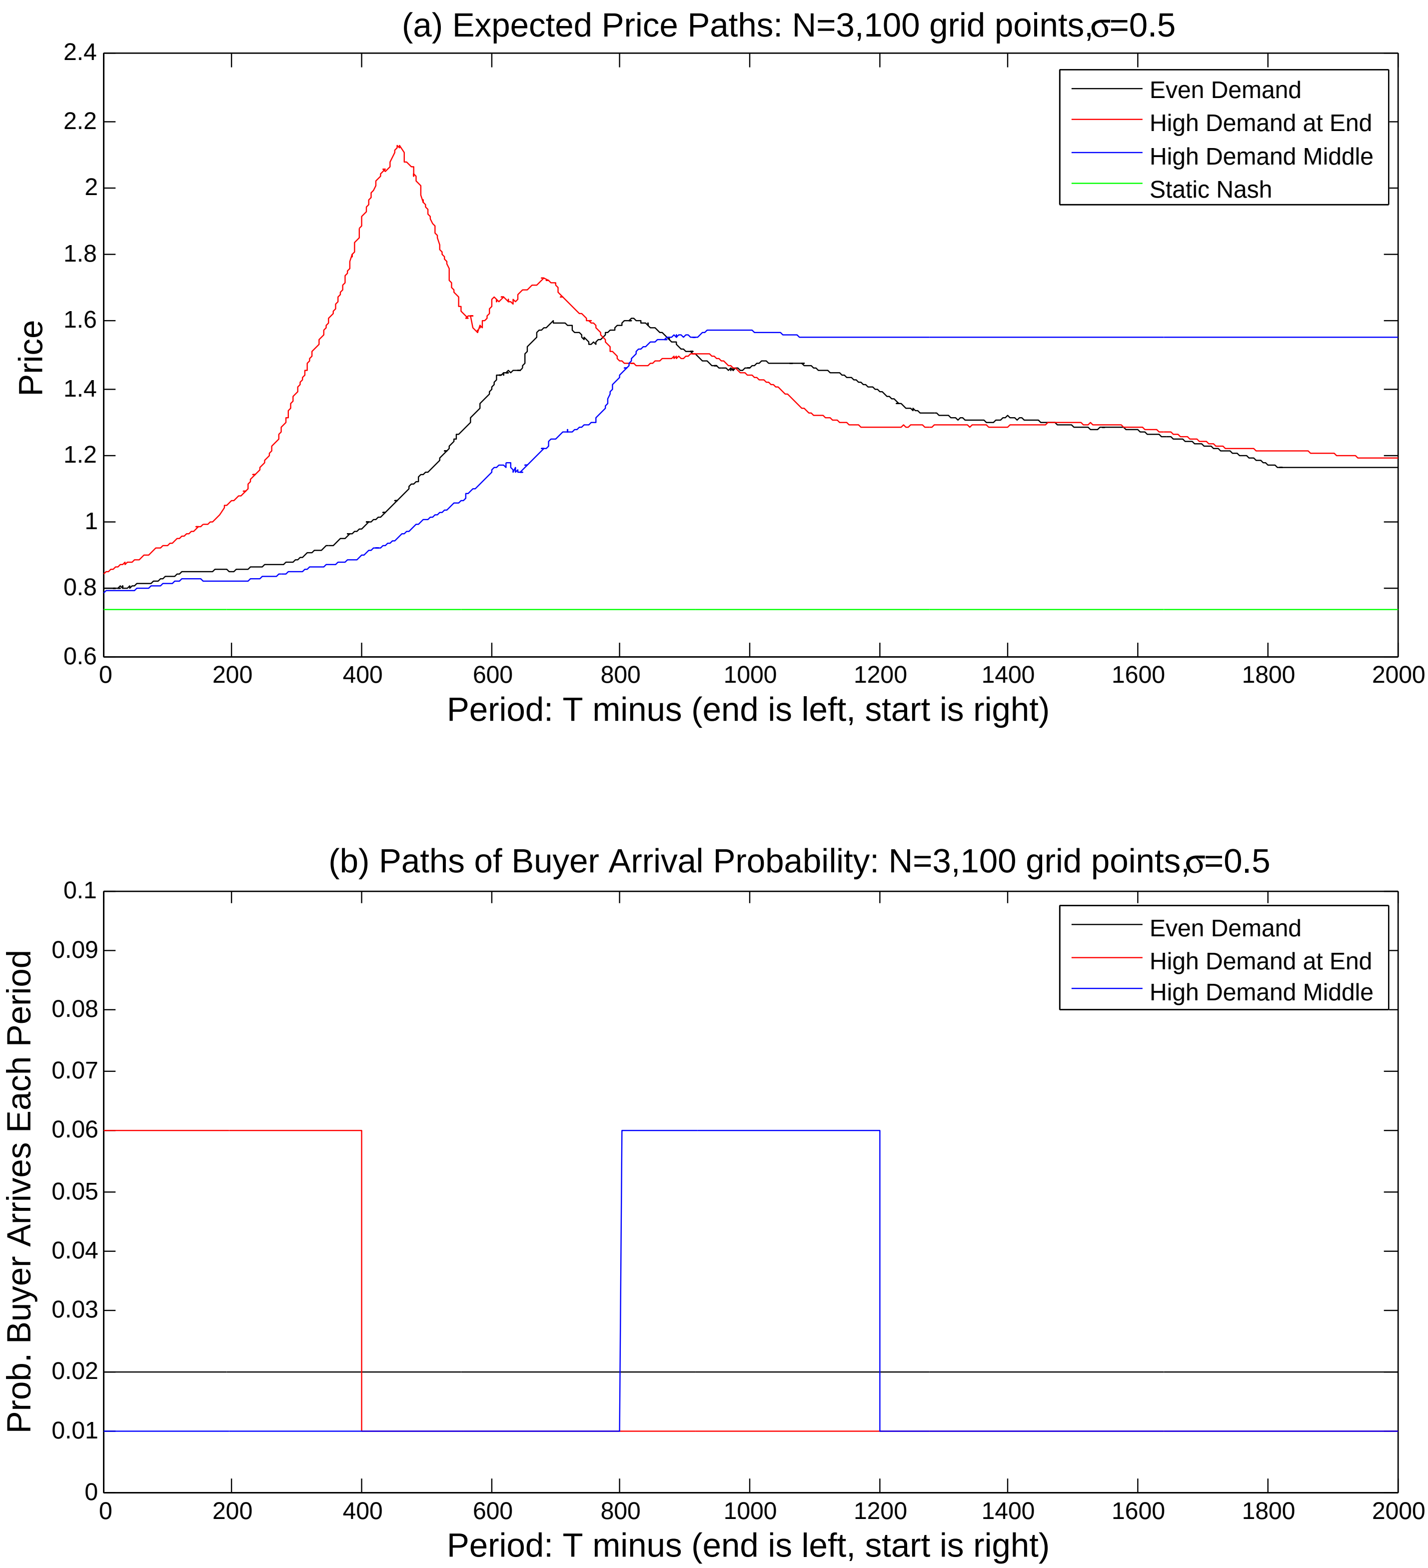

higher demand in the middle 400 periods. In both cases we lower demand in other periods so that the expected number of buyers arriving in the whole game is the same as in the base case (the path of the buyer arrival probabilities are shown in Figure 2(b)). When an increase in demand is anticipated, this increases the incentive of players to try to raise their prices in advance, if this will also tend to raise their price of competitors, as the profits from undercutting when demand is low are relatively small. On the other hand, when demand is high and prices are high there is a strong incentive to undercut; or when demand is high but expected to fall relatively soon in the future there is also an incentive to undercut to take sales now. These arguments can explain why, in the case of the blue line, equilibrium prices are higher than in the base case up to the point where demand increases, but they then start to fall at around $t = T - 950$ as the end of the high demand period approaches. When high demand is at the end of the game, expected prices increase up until period $t = T - 450$, which is later than in the other two cases, and prices reach a higher maximum level, but they then fall rapidly towards the end of the time horizon. In our empirical setting, demand is clearly much higher at the end of the time horizon, and this example suggests that this fact, combined with the non-simultaneous nature of price-setting, might explain why prices fall so dramatically as an event approaches, and why a seller, such as the broker, might be willing to set higher prices further from the event.

4 Data

4.1 Sources and Setting

Our primary dataset was provided by Stubhub.com, 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. A set of tickets sold by a seller in a particular section and row will appear as a single listing. Buyers are able to view any of the tickets that are available. For large events, hundreds or, for the largest events, such as NFL games, thousands, of listings may be available at any point in time. When a buyer first visits the buy page for a particular game they are presented with a clickable seat map - a click on a section will show them the listings available in that section - and, on the side of the page, a list of all of the listings currently available, ordered by price (lowest first) by default. The seat map makes it straightforward for buyers to compare listings in the same section even if they have different prices, and it may lead to listings in the same section being closer substitutes than those in different sections even if, ultimately, seats in several different sections might provide spectators with a very similar viewing experience. A feature of Stubhub - which makes it different from eBay - is that seller identities, or information on seller experience or performance, are not revealed. Stubhub overcomes any concerns about seller reliability by providing a guarantee that tickets at least as good as those purchased will actually be supplied, a commitment that it is able to honor because of the large number of tickets that it supplies and the fact that it maintains physical locations at or close to venues. Stubhub makes profits by charging commission fees to buyers and sellers. During our data, buyers paid commissions of 10% and most sellers paid commissions of 15% with some discounts for larger sellers.[20] 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 data contains listings and transaction data for twenty regular season home games (which we will call "fixtures" in what follows) of a US Major League sports team played between mid-January and April 2010. We focus on these games because these were identified to us by the large ticket broker mentioned in the introduction. The listings and transaction data start on January 1, so that we observe different periods of time for each fixture, but, as we will show, most transactions and active competition between sellers takes place in the last three weeks before the event, and we observe at least the final two weeks for all of the events in our sample.

The transaction data tells us exactly when each transaction took place (down to the second), the fixture, the user IDs of the buyer and seller, a numeric listing id, section and row of the tickets being purchased, the number of seats purchased and the selling price.[21] The listings data tells us the characteristics of every available listing (e.g., fixture, 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)) at any point where the characteristics of the listing changed. So, for example, there is a new record for the listing when some tickets are purchased, the price is changed or the seller adds or deletes tickets or removes the listing, possibly because he sells tickets offline. Once again we now the time of the change down to the second. We can use this information to reconstruct exactly what listings were available at any point of time.[22] While Stubhub lists tickets for luxury boxes/suites,

[20] In 2010 commissions and shipping fees were added after the buyer had chosen tickets. Subsequently Stubhub has switched to an "all-in" pricing model where the price shown on the buy page is the final price paid by the seller.

[21] In many cases we also know the seat numbers, but we do not use this information.

[22] An earlier version of this paper used data that was scraped from Stubhub's 'buy page' for each fixture every three hours for the same fixtures, so that we were able to compare the listings that a buyer was able to see with what we re-construct using the listings data, and verify that the two sets of listings matches almost exactly.

and parking passes, we completely exclude the small number (<2%) of listings of these types from our data. Stubhub does not collect data on face values, but the primary market single game ticket price for tickets in each section were added from the team's website, together with information on every team's position within its league on each day (rank).

Secondary market prices are obviously affected by the demand of fans to watch the team and what happens in the primary market. The team in our data had a successful 2009-10 season, and the reported attendances for our fixtures averaged 96% of capacity, so that, based on our conversations with the largest seller in our data, there were typically few tickets available in the primary market in the month before the game especially in the better seating areas of the venue. On the other hand, as we shall see, this does not mean that secondary market prices necessarily exceeded single game face values in the primary market.[23]

4.2 Buyers and Sellers

We model buyers as having unit demands in the sense that they will only buy once, so that they ignore how any purchase that they make will affect the future evolution of prices.[24] This is a reasonable assumption given our data: 97.4% of the 6,286 transactions we observe in our data are by buyers who only buy once for the fixture in question. 2.5% buy two listings and these are often in adjacent rows in the same section.[25] In the current version we will also assume that buyers are non-strategic, in the sense that we will follow Gallego and Van Ryzin (1994) and McAfee and Te Velde (2006) by assuming that potential buyers arrive in the market at an exogenous arrival rate and they must buy at once or depart the market forever.

While buyers are small, a feature of our data is that there is one particularly large seller on the supply-side, who we will call the 'broker' in what follows. This broker accounts for 31.5% of transactions (31.7% of seats sold) for our fixtures, and he accounts for 47.2% of transactions in the lower two tiers of the arena where face values and secondary market prices are higher. No other single seller, and we observe 344 of them, accounts for more than 6% of transactions, and only two

[23] Many transactions in the primary market take the form of whole or partial season ticket purchases, where prices may be significantly below single game face value prices.

[24] The large commission that a buyer has to pay, together with the fact that prices tend to decline as a game approaches, limits the possibility for buyers to try to arbitrage prices in this market, by speculatively buying underpriced tickets to then re-sell them.

[25] There is one instance of a buyer making nine transactions for the same game: while this buyer may affect his future prices, it is unreasonable to believe that the presence of one buyer like this, who is still small relative to the market as a whole, would significantly affect sellers' pricing behavior which is our focus.

other sellers account for more than 3% of transactions. The broker's share of listings is substantial but still smaller than its share of transactions: he accounts for 19.2% of all listings and 27.1% of listings in the lower tiers (proportions are 16.0% and 23.0% when we focus only on the last 30 days before the fixture). We will show evidence below that this seller's pricing behavior 'moves the market' in terms of the prices set by sellers who compete with him in the same section.

Equilibrium outcomes are obviously affected by sellers' objective functions. For a ticket broker this may depend on profit-sharing arrangements with whoever it is getting tickets from, and any value it gets from tickets that are not sold on Stubhub but may be sold on other online sites or offline (telephone sales or person-to-person transaction). Fortunately we were able to discuss these issues with the broker, who also provided us with its complete sales data for the fixtures in our sample. The broker clearly identified its objective as revenue maximization over the course of the event.[26] 93% of the broker's transactions took place through Stubhub, indicating that a modeling simplification where Stubhub is its only sales venue is not unreasonable.[27] We will discuss the broker's observed pricing strategy below. Of course, we know less about the objective functions, outside opportunities and sophistication, of the other sellers in the market. In this paper we will assume that they are sophisticated and make specific assumptions about their outside opportunities, and investigate the effects of the staggered price-setting conditional on these assumptions. In a companion paper we consider the optimal pricing policy of the broker allowing for the possibility that other sellers are not fully optimal.[28]

4.3 Summary Statistics

We now present some summary statistics from our sample, before discussing some stylized facts that support the idea that a seller's pricing policy can substantially affect the pricing behavior of other sellers who are listing tickets in the same section of the arena.

Figures 3(a) and (b) shows the average number of listings available at the beginning of each day[29] in the final month before a fixture, both for the broker and for all other sellers combined, and

[26] The broker acted as a sales agent for the supplier of the tickets, keeping a proportion of the revenues made, and returning any unsold tickets.

[27] Most of the remaining sales occur more than 6 weeks before the game, and so are outside the 45 day period on which we focus.

[28] For example, over 41% of sellers only have tickets in one fixture-section in our entire sample, and over 51% have less than three. They may be unaware of how prices tend to evolve and the value of particular listing characteristics, and so may simply tend to copy the prices set by sellers of similar listings.

[29] Time is measured backwards from the beginning of the fixture. So if the event start at 7pm, one day before also begins at 7pm.

Figure 3: Evolution of the Average Number of Listings Available and the Timing of Transactions

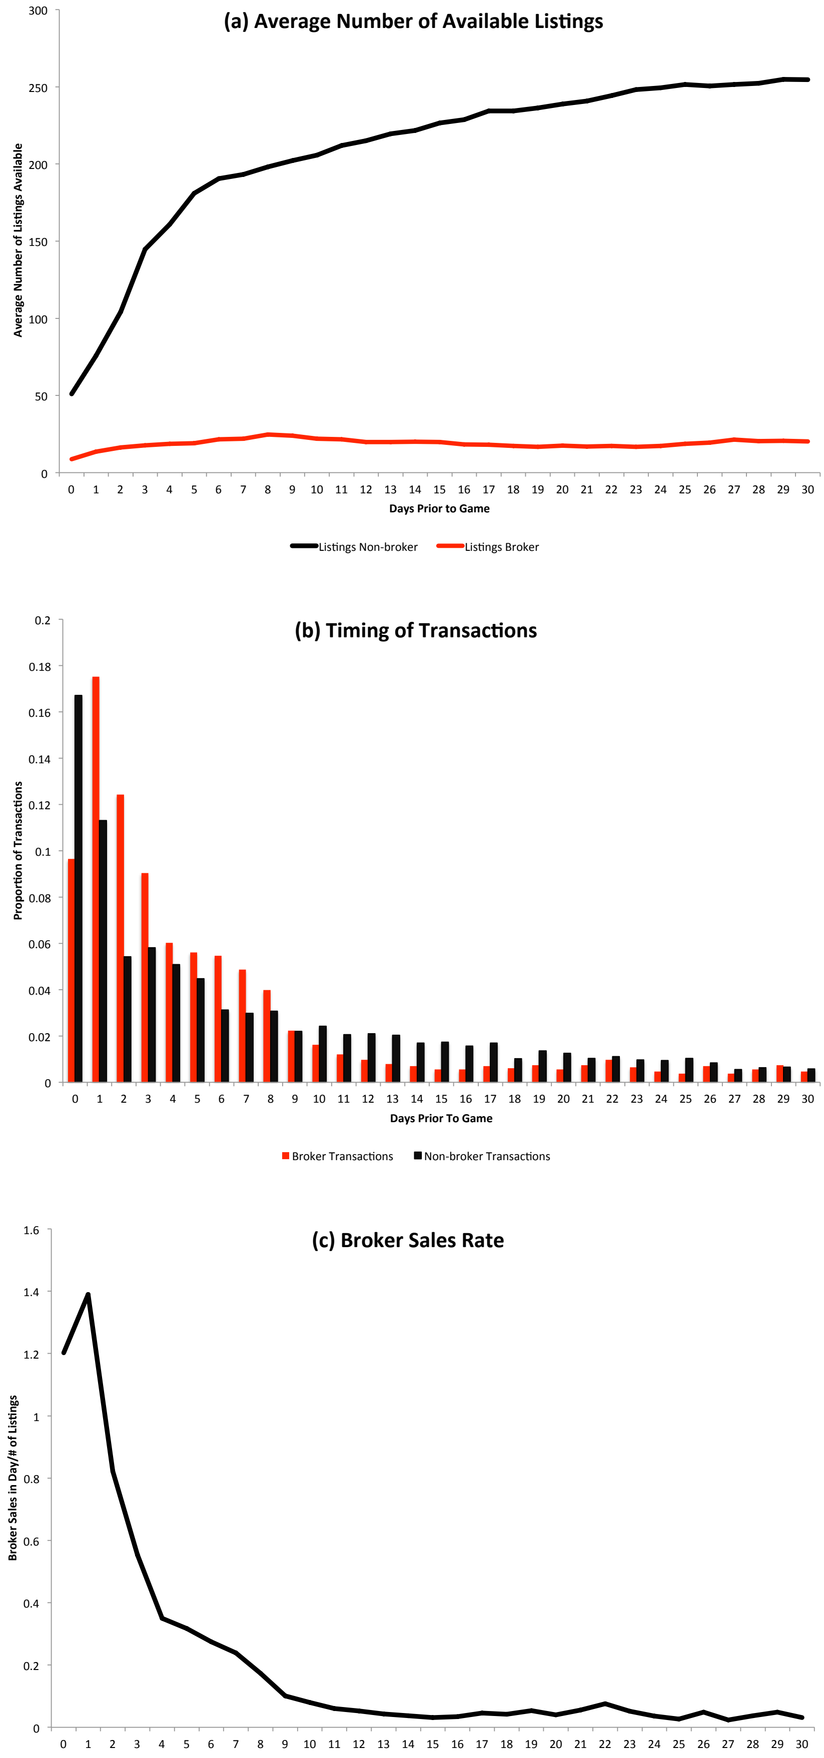

also when transactions take place. Note, that the broker's share of listings available at the start of any given day is usually lower than the 16% share reported above, because the broker is more likely to sell its tickets, so its listings are usually on Stubhub for a shorter period of time. The number of listings declines as the fixture approaches partly because of transactions, but also because tickets are delisted - either by sellers who decide to use them or sell them elsewhere - but also as Stubhub requires tickets to be de-listed because there is not enough time to ship them to the seller. When this happens depends on the type of shipping. Electronic tickets can remain available up to the day of the game, whereas paper tickets depend on the Fedex shipping options that the seller allows the buyer to select; the seller can also provide physical tickets to Stubhub which Stubhub can provide to the buyer at or close to the stadium, in which case they can remain available until game time, a format known as "Last Minute Sales"). The broker, who has electronic tickets, has an increasing share of listings as the game approaches.[30]

The number of transactions increases dramatically in the final ten days before a fixture: the median time of sale is 6 days prior to the fixture for transactions made by non-brokers and 4 days for the broker. An implication of this pattern is that when they are initially setting prices sellers should care about how they may affect price levels closer to the game. The broker accounts for the majority of transactions 24 to 72 hours before the fixture, although this share drops in the final day partly because the broker does not utilize the Last Minute Sales option. Figure 3(c) illustrates how the broker's sales rate (measured by the number of sales the broker makes divided by the number of listings that it has at the beginning of the day) increases with the approach of the fixture. It actually exceeds 100% in the last 72 hours before the fixture, reflecting the fact that a buyer may purchase only a subset of the tickets available in a particular listing.[31]

Table 2 shows some summary statistics on listing attributes, where an observation is a listing available at the start of each day before the fixture, and the attributes of listings that experience transactions (prices are those reported when the transaction takes place). The broker's seats are all in the lower and middle tiers, and, compared with non-broker seats, more of them are at the ends of the arena rather than along the sideline. The broker also lacks seats in the first two rows of each section, which are largely owned by season ticket holders, some of whom may sell them on

[30]The broker also adds some new listings in the last two weeks before some of the fixture, although many of these involve placing additional listings for sections and rows in which the broker already has a listing posted.

[31]This is possible because each buyer may purchase only a subset of the tickets that are available from a given listing.

Stubhub if they cannot attend.[32,33] The broker's listings are also more likely to have more seats than non-broker listings (4.4 to 3.5). When we aggregate across a broker's listings in the same section he has an average of 14.5 seats in each section where he has tickets, and a maximum of 101 seats in one section.[34]

Even though the broker's listings have higher face values, the broker sets lower average dollar prices for its seats than other sellers. This comparison holds true even when we exclude the small proportion of non-broker listings that have exceptionally high prices, by dropping the prices that are more than five times face value. This rule drops 1% of listings (none for the broker) and no transactions.[35] The characteristics of purchased tickets are similar to those of listed tickets[36], although, reflecting the price sensitivity of consumers, they tend to be bought at lower prices. On average, buyers purchase 2.8 seats per transaction, and 59% of purchases are of exactly two seats and 97.6% are for six or fewer seats.

4.4 Stylized Facts and Evidence of Strategic Interactions

We now describe the evolution of prices over time, and present evidence of local interactions between sellers' prices. By local interactions, we mean that the price set by one seller affect the preferred prices of other sellers who are close by in product space (which, in our case will mean the same section of the arena). While it is natural to expect local interactions in a spatially differentiated market, one might have expected that they would be limited here might be small because differentiation between sections might be limited and the number of active sellers might tend to limit the ability of any seller to charge a mark-up. We focus on the effect of the broker's prices, because the broker is a significant part of the market and his pricing policy is clearly different to other sellers.

Stylized Fact 1: sellers tend to cut prices as a fixture approaches. Table 3 shows the

[32]The broker does have one set of tickets in the front row of a section, but these tickets remain on Stubhub for less than a full day, so they are not included in these summary statistics.

[33]Row numbers are calculated by counting back from the front row of the section using a detailed seating chart (so that, for example, a row that would appear as number '5' on the ticket is counted as row '4' in our analysis if the front row in the section is numbered '2').

[34]Sellers are able to specify some quite complicated splitting rules that may allow, for example, a buyer to purchase two seats, but not one or three, from a listing with four seats. 99% of listings with four or more seats allow some number less than the total number of seats to be purchased, but many do not allow all combinations.

[35]As noted by Dinerstein, Einav, Levin, and Sundaresan (2014), a common feature of many online marketplaces is that some prices are exceptionally high (the highest price in our data is $9,000 per seat), and we should not expect our models to rationalize them.

[36]It is noticeable that purchased tickets tend to have lower face values that the population of tickets listed by the broker, while they tend to have higher face values for non-broker sellers. The latter pattern is partly driven by the fact that low face value seats in the upper tier are more likely to remain unsold.

results of using the following regression specification for the broker and all non-broker listings:

$$p_{ift} = X_{ift}\beta + D_t\alpha + FE_{if} + \varepsilon_{ift}$$

where X_{ift} are observed i (ticket) or ft (fixture-time) characteristics (e.g., the number of seats, the number of the row and measures of the performance of the home and away teams that may affect demand) and D_t are dummies measuring the number of days prior to the game. Prices (p_{it}) are the Stubhub list price divided by the face value of the ticket (the "relative" price), or the log of the list price. In the last 48 hours before a fixture, the broker has average relative prices that are under one-third of face value, and less than one-half of the average list price of other sellers.

The various columns differ in whether FE_{if} are separate fixture fixed effects and section fixed effects, or whether seller-fixture-section fixed effects, as well as what tiers of the arena are included for sellers other than the broker. When the first set of fixed effects are used the D_t coefficients indicate how the level of available prices changes over time, whereas with seller-fixture-section fixed effects the coefficients show how individual sellers tend to adjust prices if they change them.[37] Observations are the set of available listings at any point when one listing for the fixture changes its characteristics, so that the number of observations is very large and the same listing may appear in the data several hundred times.[38] To account both for this correlation, and possibly correlations across listings for the same fixture, we cluster the standard errors on the fixture.

The results indicate that as well as setting much lower prices close to the fixture, the broker also lowers its prices steadily and substantially (by 40-50% of face value) in the month before the fixture. In contrast, on average, the prices of other sellers fall much more slowly until the last four days when they drop sharply, by 15-25% of face value.

The coefficients for other ticket characteristics that are reported in the table are generally plausible. For example, seats tend to have lower prices when they are further back in a section (this pattern is especially clear for the broker). The number of seats in the listing results must be interpreted in light of the fact that potential consumers may arrive wanting different numbers of seats, so that the probability of selling one seat may actually be lower than selling two or four seats even though in a classic dynamic pricing model where all potential consumers have unit demands,

[37]If they have multiple listings in a section which are created at different times, the coefficients can also capture how the relative price of these listings varies with the date that they are posted. Cross-listing variation also identifies coefficients on the number of seats.

[38]Results are similar if one uses the first observation from each day.

optimal prices will tend to be higher when there is a single unit to sell. Surprisingly the league rank results indicate that a lower rank for the home team tends to increase the broker's prices, while team rank has very little effect on non-broker prices, especially in the lower tiers. This may be explained by the fact that the team in question had a consistently successful regular season where slightly worse performance made upcoming games more potentially pivotal for getting through to the post-season.[39]

Stylized Fact 2: the broker tends to cut prices more dramatically than other sellers, and its pricing policy is relatively insensitive to the number of seats that it has left.

The regressions in Table 3 illustrated how the broker tends to lower prices more dramatically as a game approaches than other sellers. This pattern is somewhat surprising given that the broker has more inventory to sell and monopoly models of dynamic pricing tend to predict that multi-unit sellers should tend to adopt relatively flat pricing schedules. We can also test an additional prediction of these models, which is that optimal prices should tend to increase when the seller has less inventory to sell, whereas they should be lower when there is more inventory to sell.[40]

Table 4 repeats the regressions for the broker, using relative prices, but including the total number of seats that the broker has in the section and this variable interacted with the log of the number of days until the fixture as additional explanatory variables (the mean of this variable is 20 and the standard deviation is 33). We also estimate specification where we include the square of the number of tickets remaining. The results tend to indicate that the broker sets lower prices when it has more units to sell but that the effects of differences in inventory are fairly small, especially further from the game. For example, the coefficients in column (2) indicate that the broker increases its prices by an insignificant 0.5% of face value (s.e. 1.3%) when it has 30, rather than 20 tickets to sell, 25 days before the fixture; and this same increase in inventory would cause the broker to only lower its prices by 2.6% of face value two days before the fixture. While one would obviously need to completely specify a model to figure out how sensitive a monopolistic dynamic pricing model would imply that prices should be to different levels of inventory, one possible explanation for why the effects are so small is that, if the broker's prices affect the prices that other sellers set, it will care more about raising its competitors prices when it has more prices left to sell, and this may

[39] As the broker changes its prices more regularly than non-brokers, especially before the last few days, the broker's pricing may be more sensitive to slight changes in the likely degree of interest in the game.

[40] This outcome can also happen in models with competition as a seller may try to aggressively sell off inventory when it has a large amount in order to try to increase the future prices that it will want to set, as this will strategically raise the prices of its competitors.

offset the usual incentive to lower prices.

We were able to discuss the broker's pricing strategy with the broker. At the time of our data the broker did not have an automated system for updating its prices. Instead it used a team of employees to change its prices as required. While these employees had discretion about setting any particular price, the broker's manager explained that his agents usually set prices that were lower, or close to being lower, than other sellers in the section, especially as the fixture drew near, but that, at the same time, they were reluctant to undercut other sellers too much on the basis that they would lower their prices in response.

To gauge whether this is an accurate description of current pricing, Figure 4 plots how the average proportion of the broker's listings that have the lowest prices in their sections changes over time. The figure is based on the set of listings available at the start of each day. Obviously this proportion may be mechanically affected by the number of listings available in the section changes (in particular, it will tend to be higher when there are fewer non-broker listings). Therefore the figure also shows, as a comparison, the probability that a broker's listing would be the cheapest if a listing in the section was randomly selected to be the cheapest.[41]

The figure broadly confirms the broker's description of his strategy, as, on average, the broker line lies above the comparison line. It is also noticeable, however, that this pattern is less marked 12-20 days before a fixture. In Table 3 one can also see that the decline in the broker's prices is broken by an increase in prices around 17-19 days to go, so that, on average, prices are at the same level or slightly higher 11-13 days before the fixture than 20-24 days before the fixture.[42] Partly because of its behavior during this period, only 69.1% of the broker's price changes are price reductions, compared with 88.3% of price reductions for non-brokers even though the size of the total decline in in the broker's prices in the last month before a game is larger. One interpretation of this pattern is that the broker, who knows that demand will be much higher in the last two weeks before the event, holds off cutting its prices in order to try to raise his competitors' prices (or at least prevent them from falling as much as they otherwise would) as this increase in demand approaches.

[41] To illustrate how the comparison probability is calculated, suppose that at the start of the day there are three broker listings and seven non-broker listings in a section for a particular fixture. An indicator for the broker's listing being cheapest is equal to one if, out of the ten listings, one of the broker's three listings has the lowest price. The comparison probability based on a random selection would be $\frac{3}{10}$. In the figure the average is taken across fixture-sections where the broker has listings and there is at least one non-broker listing. While the comparison probability rises, the patterns are similar if one excludes non-broker listings in the first two rows of a section.

[42] To illustrate using another metric, 78% of the broker's price changes 19 days before the fixture are price *increases*. As a comparison, 8 to 14 days before a fixture no more than 17% of price changes are price increases.

Figure 4: Broker's Pricing Strategy: Proportion of Sections Where the Broker Has the Lowest Price

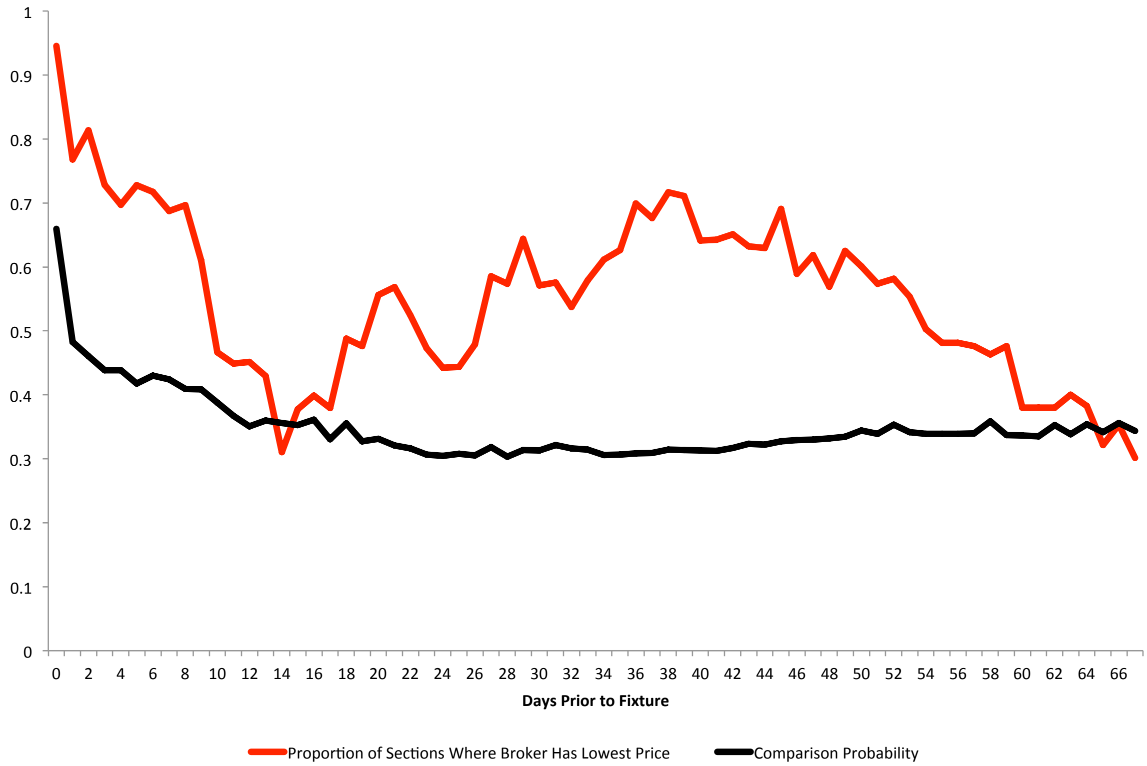

Stylized Fact 3. Competition from the broker affects the pricing of other sellers in the same section. A key assumption in our model is that the broker's prices may influence the prices that other sellers set. While the fact that the broker does not drop its prices significantly when it has more units to sell, and the fact that the broker seems to raise its prices just as the market activity tends to increase are suggestive that strategic incentives affect the broker's pricing decisions, we can also assess the effect that the broker's prices have on other sellers more directly.

We identify three related pieces of evidence. First, when we look at the last four days before the game, which is when more than 40% of transactions occur, other sellers set lower prices when the broker is listing tickets in the same section. To show this we estimate the following specification using the sample of *non-broker* listings

$$p_{ift} = X_{ift}\beta + FE_{if} + \varepsilon_{ift}$$

where p_{ift} is the relative or log price of a non-broker listing, X includes characteristics of the listing (row, number of seats in the listing, team form) plus a dummy variable for whether this is a game-section where the broker has seats.[43] FE_f are fixture fixed effects.[44] Differences between sections are captured by either including section fixed effects or including a set of variables that identify particular areas of the arena (lower tier, middle tier, a premium seating area in the middle tier, and dummies for each of the ends). When we include section fixed effects, identification of the broker competition coefficients comes from cross-event variation in whether the broker has seats in a particular section; whereas when we only include section characteristics we also get some identification from within-fixture variation across similar sections in whether the broker is present. The coefficients reported in Table 5 indicate quite substantial competition effects: non-broker prices are 10-20% of face value lower when the broker has tickets in the section. One obvious concern is that this pattern might reflect the presence of additional tickets rather than the broker's lower prices. Therefore in the final three columns of Table 5 we repeat the relative price specifications including variables that measure the total number of listings and the total number of seats available in all other listings in the section as additional controls. The coefficients on these controls are very small, and coefficients on the broker presence dummy variable are scarcely affected by their inclusion (this is also true if we include higher order terms for the number of listings and tickets).

The second piece of evidence is that in the presence of the broker, other sellers cut prices more quickly, as well as setting lower prices, in the month leading up to the game. We assess this by repeating the price path regressions (Table 3) using non-broker listings with a couple of changes. The first change is that we add a set of interactions between the days-prior-to-fixture dummies and a dummy that is equal to one if the broker does not have listings in the section, defined either as whether the broker never has listings in the section for the fixture in question (cols. (1)-(3), "ever") or whether the broker currently does not have listings (cols. (4)-(6), "currently"). The second change is that we now use listings more than 30 days prior to the game as the excluded group, making it easier to compare the levels of prices at the beginning of our time period and how they decline from this point. All of our specifications include fixture fixed effects, and we either include section fixed effects or section characteristics. The results in Table 6 show that non-broker prices tend to be higher one month before the fixture when the broker does not have tickets in the section

[43] The variable is equal to one if the broker ever lists seats in the section for the fixture in question. Results are also significant using dummies defined as being equal to one if the broker currently has seats in the section or has seats in the section 24 days before the fixture. We show some specifications using current competition from broker listings in later sections. When all of these variables are included, it is usually the case that the "ever" dummy is significant and the others are insignificant, although this result does depend on the exact specification that is chosen.

[44] Observations are the set of listings that are available whenever any listing for the game changes one of its characteristics. Therefore, as before, the same listing can enter the sample many times, which we correct for by clustering the standard errors on the fixture.

(the coefficient on the Broker Does Not Have Tickets dummy), and that from this point until two days before the game prices decline more slowly in these sections than in sections where the broker does have listings (the positive and significant interactions between this dummy and the days-to-go variables).

The final piece of evidence comes from looking at how the presence of the broker affects the frequency of price changes and price reductions. Figure 5 plots the proportion of listings that experience price changes and price reductions each day in the last month prior to a fixture, distinguishing between the broker's listings, non-broker listings in fixture-sections where the broker (currently) has listings and those where it does not have listings.[45] We restrict analysis to the middle and lower tiers to make broker and non-broker behavior more comparable. The results show that non-brokers change, and especially lower, prices more often when the broker has listings in the same section, until two days before the game when the pattern reverses. Notice, however, that price changes by the broker do not immediately result in a spike of price changes by non-brokers. This affects how we set up our model; in particular, it would not be right to assume that non-brokers are constantly monitoring the market so that they change prices immediately.

The figures also illustrate how the broker changes its prices more frequently than non-brokers until the week before the fixture. Table 7 shows that the effect of broker presence on non-broker price changes is robust to adding controls for section and listing characteristics, where we use probits to assess the effect of the presence of the broker on the probability that a non-broker listing experiences a price reduction, price increase or price change in a given day.[46] An observation is a listing-day, and we use observations from the last month before the game. Standard errors are again clustered on the fixture. The effect of broker presence is allowed to vary with whether there are 0-2, 3-10 or more than 11 days left before the fixture. Consistent with 5, the coefficients show that broker presence is associated with significant increases in the probability of both non-broker price changes and non-broker price reductions more than 3 days prior to the game.[47] Broker presence reduces the

[45]We first identify how many price increases and price reductions that a listing has during a day (typically there are no changes or one change, but sometimes there are more than one). If a listing has more reductions than increases, then this counts as a price reduction for the day, and the opposite if there are more price increases. A listing has a price change if it experiences at least one price increase or price reduction, even if there are two changes that leave the price the same.

[46]Note that we are estimating separate probits for each type of price change, and do not account for the natural dependency that results across the specifications.

[47]The sum of the broker presence coefficient and the coefficient on the interaction and the 3-10 days-prior-to-fixture dummy is statistically significant at the 5% level in all of the columns involving price reductions or price changes. The sum of the broker presence coefficient and the coefficient on the interaction and the more than 11 days-prior-to-fixture dummy is statistically significant at the 5% level, apart from in columns (10) and (11) where they are

Figure 5: Time Path of Non-Broker Price Changes and Price Reductions as a Function of Broker Competition

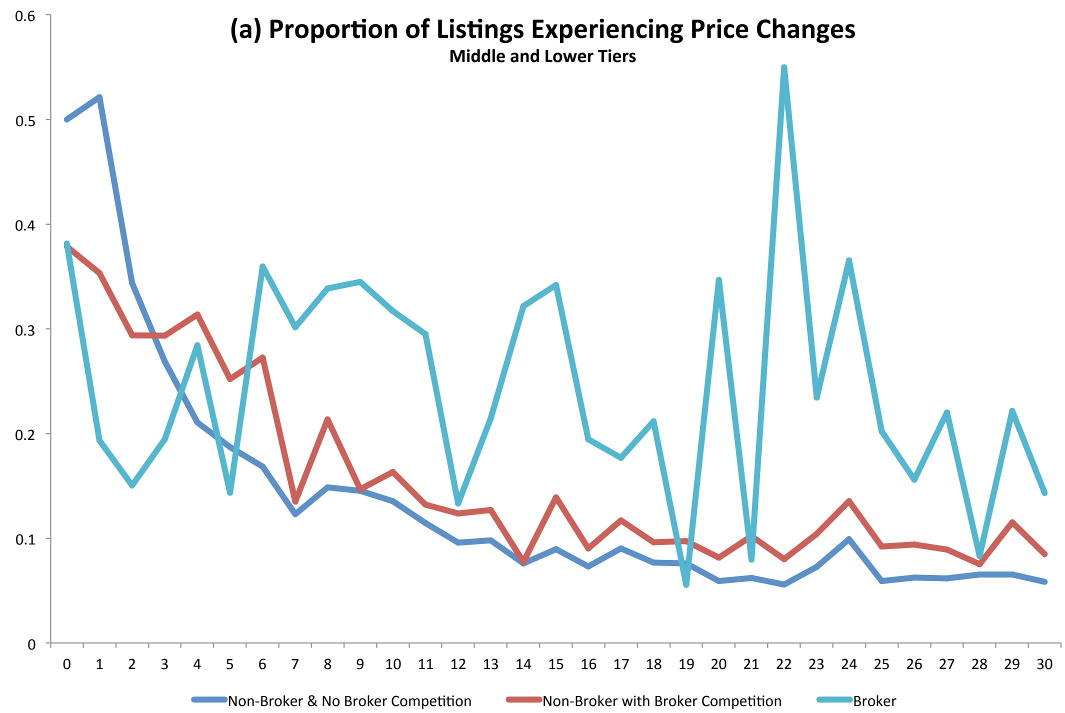

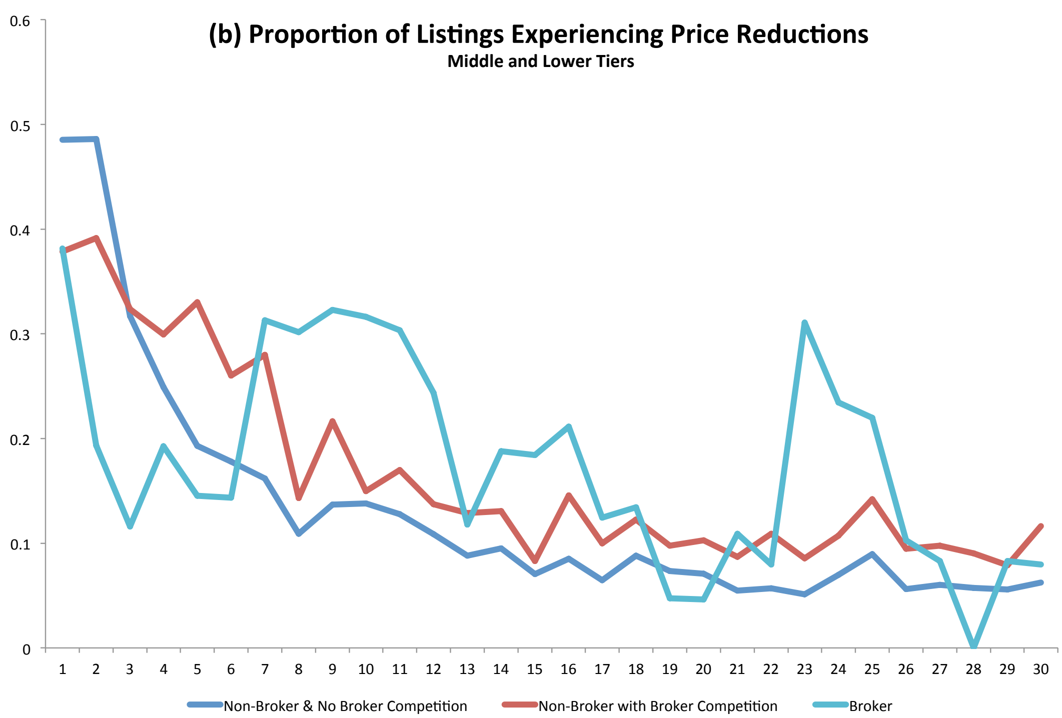

probability of non-broker price increases, although these are rare events anyway and the effects are not statistically significant.

Before we turn to the model, we note that we have not explicitly controlled for the shipping options associated with each listing (electronic, Fedex, available for last minute pick-up from Stubhub close to the arena). These are relevant in that shipping options may be valued differently for different buyers and, in 2010, the shipping option affected the total price that a buyer had to pay (so far all of the prices used have been the price of tickets without either buyer commissions, which should be the same proportion of the prices that we have used for all listings, or shipping fees, which could vary between $4.95 for electronic tickets to as much as $25 for the highest priority Fedex option, although these prices are per transaction rather than per seat). This reflects one limitation of our data where the shipping options associated with each listing are not systematically reported, although we do observe the selected shipping option for transactions. For the broker, all but two of its Stubhub transactions involve electronic tickets that have the lowest shipping cost. For other sellers, 15% of transactions are electronic, 30% involve last-minute pick up and the remainder involve some form of Fedex shipping. Therefore, while it would obviously be preferable to control for shipping options, including the cost of shipping would certainly not change our observation that the broker sets significantly lower prices than non-brokers (in fact, it would tend to strengthen this conclusion).

5 Model

The model that we set-up in order to measure the role that various factors play in causing prices to fall as a game approaches incorporates both sequential price-setting and product perishability, as well as allowing for a large seller who may have different incentives that small sellers. That said, it is obviously necessary to simplify along some dimensions, especially as our experience with sequential price-setting models in Section 3 indicated that the size and richness of the price grid can matter for the level and dynamics of equilibrium prices when prices are set sequentially (in particular, average prices and the size of the price declines at the end of the game tend to be larger when a finer grid is used).

We focus on competition at the level of a section within the arena, modeling competition for tickets in other sections and the profits that the broker can make from making sales in other sections

insignificant at any conventional significance level.

as evolving according to exogenous and deterministic processes. In our representative section, there is, at most, one 'large' seller (b) who can have at most n_b units of inventory. There can be at most N_s other (small, s) sellers with at most n_s units each. From this point onwards, it is useful to think about a unit as being a pair of tickets. The game is one of complete information, so that every seller knows the identity and inventory of all other sellers. This is realistic in the sense that the broker does not strategically withhold inventory but it is unrealistic in the sense that listings are anonymous so that the type of owner for each listing is not public information.[48]

There are T (short) time periods. At the end of the game, each unit of remaining inventory of a small seller is worth v_T^S. The large seller has no value for unsold inventory so that its aim during the game is to maximize revenues. The state space at the start of a period t consists of the inventory and current prices of each seller, the period index, and the value of certain additional state variables that will be assumed to evolve exogenously. In each period, nature selects a particular type of event to occur (or none at all).

Buyer arrival. A potential buyer arrives with probability $q^{\text{buyer}}(t)$. The buyer chooses whether to buy a ticket in the section of interest, or in another section of the stadium, where the choice probabilities are determined by a nested logit model where sections are nests. For the current estimates, we assume that all potential buyers only want to purchase a single unit, so that all listings are included in the choice set. However, it is not difficult to extend the model to allow for there to be buyers who only want to buy two or more units, so that the choice set will be more restricted. We could allow for an outside of not-buying from any section; however, with our current data estimating the role of the outside good would depend on functional form as we do not observe buyers arriving who do not purchase anything.[49] The probability that a unit is purchased from a

[48]This could potentially have a significant effect on prices. If small sellers make inferences about how prices will evolve based on current price levels, this might provide an additional reason why a large seller might want to set higher prices earlier on in the game. However, incorporating asymmetric information within a dynamic game is not straightforward (Fershtman and Pakes (2012) and Gedge, Roberts, and Sweeting (2014) for recent examples where asymmetric information is allowed for).

[49]Relying on functional form, one finds estimates where any arriving buyer purchases with probability very close to one fits the data best. We hope to add data that will allow us to see how many buyers arrive without buying in a later iteration of the paper.

buyer j of type $\tau(j)$ in the section of interest if a buyer shows up, is, following Berry (1994),

$$\Pr(\text{buy}_{j,\tau}) = \frac{\exp\left(\frac{\delta_{\tau(j)}-\alpha p_j}{1-\sigma}\right)}{\left\{\begin{array}{c}\left(\sum_{k\in\text{sect of interest}}\exp\left(\frac{\delta_{\tau(k)}-\alpha p_k}{1-\sigma}\right)\right)^{\sigma} \\ *\left[\left(\sum_{k\in\text{sect of interest}}\exp\left(\frac{\delta_{\tau(k)}-\alpha p_k}{1-\sigma}\right)\right)^{1-\sigma}+D_{\text{other\_sect}}(t)\right]\end{array}\right\}}$$

where σ is the nesting parameter. The $\delta_{\tau(j)}$ terms allow for listings offered by different sellers to differ in their attractiveness to buyers. While we do not explicitly model differences in competition that arise from listings having different row and shipping characteristics, this type specific difference allows for the listings of one type of seller to be more attractive on average. $D_{\text{other\_sect}}(t) \equiv \sum_{s\notin\text{sect of interest}}\left(\sum_{k\in\text{sect } s}\exp\left(\frac{\delta_{\tau(k)}-\alpha p_k}{1-\sigma}\right)\right)^{1-\sigma}$ measures the attractiveness of tickets from other sections at time t, which we calculate given our estimated demand system.

If a seller sells a listing then it receives the sale price as a payoff, although its value will also adjust to reflect the fact that after the sale it will have fewer listings left. For the large seller we also seek to account for the fact that it may receive a payoff when one of its listings in other sections sell. We allow for this by calculating, in our data and using our estimated demand, the broker's expected revenue when a sale is made if there is no sale in the specific section that we are looking at. In this way we can account for the fact that the broker will have listings in many other sections as well as the one under consideration. This payoff, $R_b(t)$, can vary over time. On the other hand, we are not capturing the fact that when a seller is setting its price in the section of interest it may take into account that its choice may affect the prices that are set in other sections.

Price change opportunity. With probability $q^{\text{p change}}(\tau, t)$ an opportunity to change its price arises for player of type τ. Only one player will get this opportunity in a period t. In this case, the player's optimal policy will maximize its value

$$p^*_{j,\tau(j),t} = \max_{p'_{j,\tau(j),t}} V_{\tau,t+1}(p'_{j,\tau(j),t}, \mathbf{p}_{-j,t}, \mathbf{n}_b, \mathbf{n}_s, \mathbf{H}_t) \tag{6}$$

where the bold variables indicate vectors and $\mathbf{H}_t$ gathers other state variables such as $D_{\text{other\_sect}}(t)$ and $R_b(t)$ that are assumed exogenous. Note that we assume that there is no cost to changing prices, and that there is no 'shock', say to marginal costs, that might cause sellers in the same ex-ante state, and of the same type, to choose different prices.

Seller arrival and exit. In our data, new sellers arrive and some existing sellers leave without selling out. The number of tickets the broker has may also increase over time.[50]

With probability $q_s^{\text{arrival}}(t)$, a new small seller arrives (if there is room in the state space). On arrival, a small seller gets to an opportunity to set an optimal price given the state of the market (i.e., an arrival is coupled with a price-setting opportunity for the arriving seller). Arrival is independent of the current state (as long as there is 'space' given the definition of the state space).

With probability $q_s^{\text{depart}}(t)$ each of the existing small sellers leaves the market. Obviously we do not know what price a small seller receives if he sells his tickets elsewhere (or what utility value the seller gets if he decides to go to the game himself). We therefore assume that he receives a payoff that is equal to his current value of remaining posted on Stubhub; i.e., he receives a payoff that is equal to his value of being in the same market-state at the beginning of the next period. This assumption means that the possibility of exit does not have a direct effect on a seller's pricing decision (he should not increase or lower his price to change the probability of exit), although the possibility that other sellers may exit will obviously affect pricing as it affects expected future competition.

With probability $q_b^{\text{add}}(t)$ the number of listings in the broker's inventory increases by one unit. In practice, the broker's inventory often increases by a number of units at once, and more generally we can think of using this probability as an ad-hoc way of capturing the fact that broker will usually expect his inventory to be greater than n_b, even though n_b is the maximum number of units of inventory that we can practically allow in the state space.

6 Estimation

To perform the counterfactuals we need to estimate demand and the main arrival rates from the observed Stubhub data. To do so, we must organize the data (which lists the times of different events to the second) into time periods that are long enough to allow the practical application of backwards induction to solve a game that lasts for 45 days before a game takes place, yet small enough that our assumption that only one event occurs during a time interval is a plausible approximation.

We achieve this in two stages. In the first-stage we discretize the data into short five minute time periods, and map the continuous time events (price changes, purchases, seller arrivals and

[50] One could assume that he anticipates that more tickets will arrive before he receives and lists them. If he is certain to get them then they should affect his listing decisions in a similar way to existing inventory.

departures) into these time periods.[51] If more than one event happens during a five minute time interval then we push the event that occurred first back into the previous time interval. This can, of course, knock some other 'events' backwards. However, when we look at individual fixture-sections this process results in no event being pushed more than about one hour from when it actually took place. We then use this five-minute data to estimate the arrival rates required for the counterfactual.

However, as we will show, the probability that any event occurs within a five minute interval is low once we are more than 10 days or so from a fixture. Therefore, we aggregate five minute time intervals where the probability that any event occurs is small so that the probability that more than one event would happen during our time aggregated intervals is less than 0.005 (we do not attempt to disaggregate five minute time intervals where the probability is greater than 0.005). As a result, the time intervals close to the fixture are short and the time periods several weeks before the event are much longer. For example there are at least 144 time periods on each of the last five days before the fixture, and 96 on the seventh day before the fixture and 35 periods twenty-one days before the fixture. We will solve counterfactuals for the last 45 days before the fixture, with 2,329 time periods in total.

6.1 Demand Estimation

As noted above, we assume that all consumers who visit will purchase at least one listing. We estimate the demand for listings using a nested logit model and all transactions, allowing the choice set to vary with the number of seats purchased and the rules associated with each listing that determine which subsets of tickets can be purchased. Sections are nests. Table 8 reports the results from three specifications. In the first specification we estimate the effect of prices (measured relative to face value), the position of the section within the arena and the row. The coefficient on price is large, negative and highly statistically significant. The only coefficient on the seat characteristics that does not make intuitive sense is the 'Premium' seating coefficient. Premium seats are located in the middle tier and give the seatholder access to additional facilities such as in-seat waiter service. However, these seats have very high face values ($125 per seat, only $3 cheaper than seats right in the front and in the middle of the arena), and the negative coefficient may reflect the fact that the type of supporter who buys tickets on Stubhub has lower demand for these features (or is

[51] We have also experimented with one minute time periods and the estimates of arrival rates are very similar.

less aware of them) than season ticket-holders or people who buy directly from premium ticket brokers). In the second specification, we allow for broker listings to have a different quality effect. While consumers do not see that the tickets are being sold by the broker this is intended to proxy for the fact that the broker's tickets tend to have more attractive shipping options than those of other sellers, and the coefficient on broker is positive and highly significant. In the third column we allow for the price coefficients to vary with the number of days until the fixture (measured as a continuous variable). The number of days to go has only a small effect on the price coefficients, which is broadly consistent with the results presented in Sweeting (2012), where he estimates a demand curve for individual listings on eBay and finds that the slope coefficients are fairly constant as a game approaches.[52] In the counterfactuals, we will use the estimates from column (2). At the bottom of the table we report the average value of the mean utility of listings excluding price effects, for game-section observations where the broker has tickets. Here we see that even though the broker has almost no seats in the first two rows of any section, the coefficient on 'broker' is sufficiently large that broker listings have higher average quality. In our counterfactuals we assume that $\delta_b = 0.377$ and $\delta_s = -0.056$ (recall, that counter the usual interpretation of mean utilities, these parameters exclude the effect of prices).

We also use the demand estimates, and observed prices, to calculate values of $D_{\text{other\_sect}}(t)$, which measures the attractiveness of listings in other sections, and $R_b(t)$, which measures the broker's expected revenues when a consumer arrives but no tickets in the section of interest are sold. We calculate these values for each of the sections where we observe the broker having tickets at the points where some sale is made, and then estimate how the average values of these variables evolve as a flexible function of the time until the fixture. The estimated time paths of these variables are shown in Figure 6. For both measures the pattern is non-monotonic, reflecting the fact that both the number of listings is declining over time and prices are falling, which can tend to push these variables in different directions.

6.2 Arrival Rates

Given our assumptions that all consumers that arrive on the website make a purchase, we can estimate the probability that a consumer arrives during a five minute time period using only observed arrivals. We allow the arrival rate to vary as a flexible (polynomial) function of the time until the

[52]Sweeting (2012) addresses the possibility that prices are endogenous by using instruments, reflecting the willingness of the seller to sell tickets, as instruments.

Figure 6: Time Path of $R_b(t)$ and $D_{\text{other\_sect}}(t)$

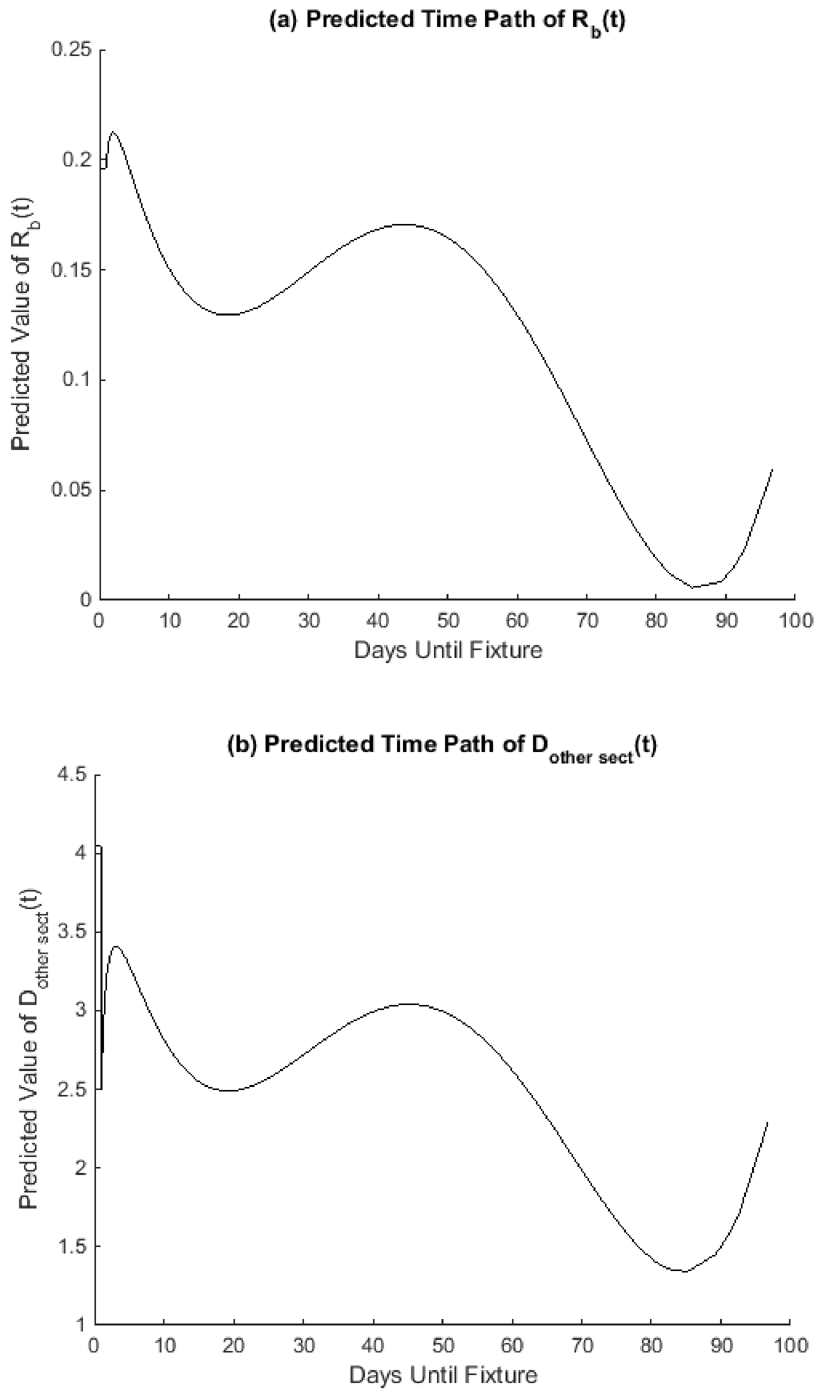

fixture, with the estimated time path shown in Figure 7(a).[53]

In the current draft we estimate the arrival of a price change opportunity for the broker and small sellers using observed price changes (i.e., we will underestimate the arrival rate by ignoring the possibility that a seller considers changing its price but chooses not to do so). We identify a price change when the seller changes the price of any of the listings that it has within a section. Figure 7(b) shows the estimated probability of price change opportunity within a five minute interval as a function of the time until the fixture, for both types of seller. Consistent with the descriptive statistics, the probability of a price change opportunity increases as a fixture approaches for both types of seller, although the probability that the broker changes its price is always higher. It is unclear why the probability that broker changes its price is higher 60 days before the game than it is 40 days before the fixture. However, in the counterfactuals we focus on the last 45 days before the fixture, so this non-obvious pattern should not affect our results.

We also estimate the probability that a new small seller arrives in the market, and the probability that a small seller exists the market as a function of the time until the fixture. The estimated time paths are shown in Figure 7(c). Both rates increase as the fixture approaches, with the probability of departure increase dramatically in the last 10 days before the game (partially reflecting the fact that listings must come off when there is no time to ship tickets to a buyer). The increasing arrival rate presumably reflects the fact that some people who planned to attend find out that they are unable to go in the last few days and try to sell their tickets.

7 Counterfactuals

[to be presented at seminar]

8 Conclusion

This paper attempts to quantify the role that non-simultaneous price setting plays in the dynamics of prices in an online market, in a setting where a large seller cares about how other sellers will respond to the prices that he sets. In doing so, we have two primary motivations. First, we want to

[53]One can also allow all of the arrival rates to vary with time of day. These results show sensible patterns: there is less activity during early morning hours, the broker is much more likely to change its prices during work hours, while buyers and small sellers are relatively more active in the evening.

Figure 7: Arrival Rates

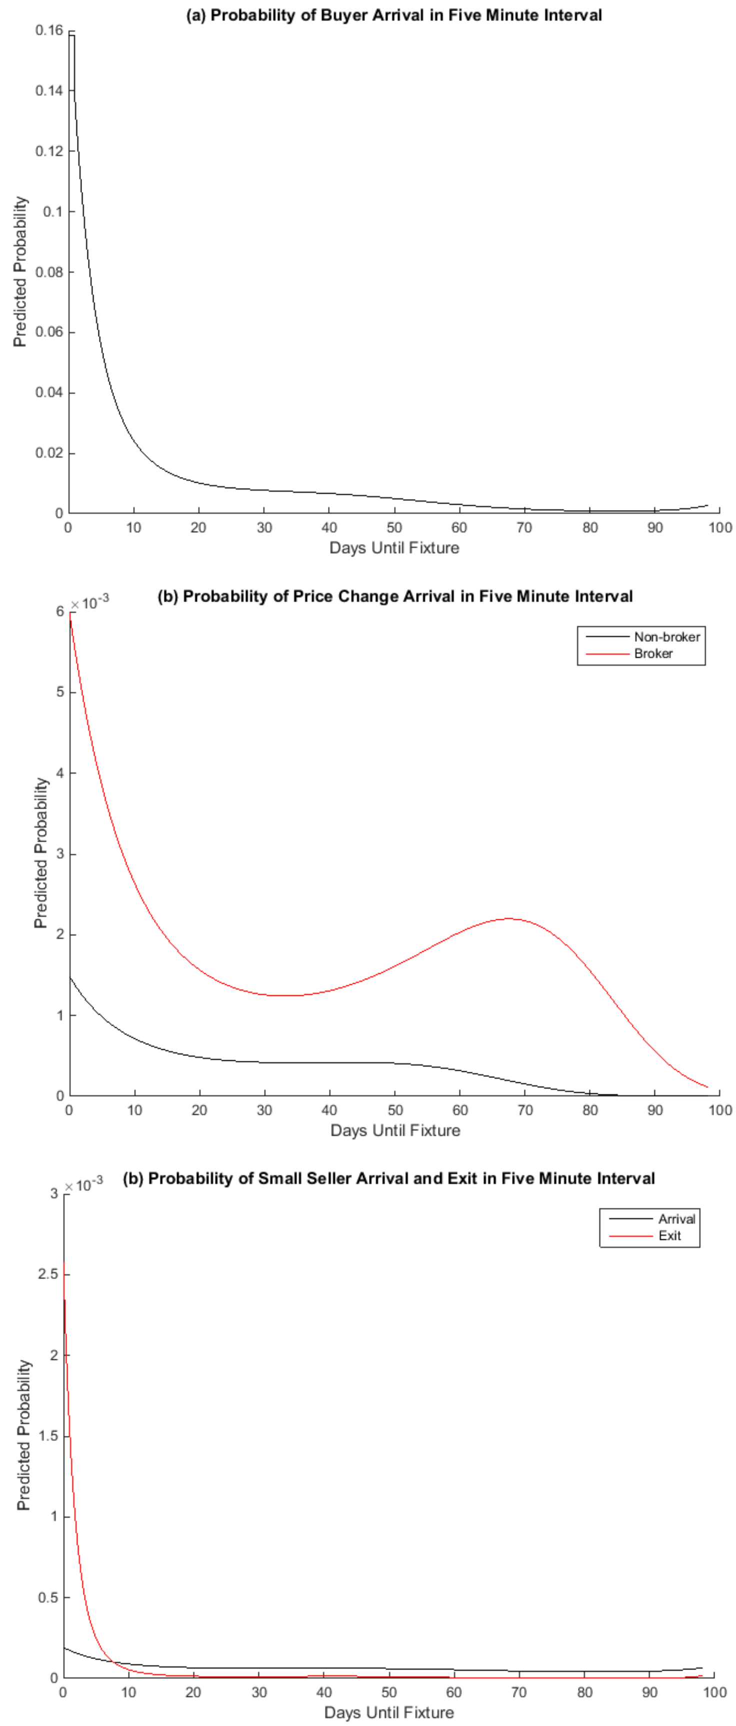

understand what constitutes a good model for understanding price-setting behavior in our type of setting and which could be used by a platform to understand elements of optimal platform design. For example, identifying whether it is optimal to offer commission discounts to very large sellers, such as the broker, or whether discounts should be given to encourage trade further before the game when prices are high, depends on understanding how any discounts will be passed through to the prices that all sellers set on the platform, and doing this requires having a good model of equilibrium price-setting behavior.

Second, our analysis constitutes a first pass at understanding the role of non-sequential price setting much more generally. As the analysis of the example in Section 3 indicates, non-sequential price setting can lead to price levels, and differences in the comparative statics of how price-levels vary with market structure, that are quite different from the simultaneous price-setting case.

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.

Arcidiacono, P., P. Bayer, F. A. Bugni, and J. James (2012): "Approximating high-dimensional dynamic models: Sieve value function iteration," Discussion paper, National Bureau of Economic Research.

Berry, S. (1994): "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, pp. 242–262.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of monetary Economics*, 12(3), 383–398.

Caplin, A. S., and D. F. Spulber (1987): "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, pp. 703–726.

Carlton, D. W. (1989): *Handbook of industrial organization*vol. 1, chap. The theory and the facts of how markets clear: is industrial organization valuable for understanding macroeconomics?, pp. 909–946. Elsevier.

De Roos, N., and H. Katayama (2013): "Gasoline price cycles under discrete time pricing," *Economic Record*, 89(285), 175–193.

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 K. L. Judd (2012): "Avoiding the curse of dimensionality in dynamic stochastic games," *Quantitative Economics*, 3(1), 53–93.

Eckert, A. (2013): "Empirical studies of gasoline retailing: A guide to the literature," *Journal of Economic Surveys*, 27(1), 140–166.

Eckert, A., and D. S. West (2004): "Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations*," *Journal of Law and Economics*, 47(1), 245–273.

Ellison, S. F., and C. M. Snyder (2014): "An Empirical Study of Pricing Strategies in an Online Market with High-Frequency Price Information," .

Ericson, R., and A. Pakes (1995): "Markov-perfect industry dynamics: A framework for empirical work," *The Review of Economic Studies*, 62(1), 53–82.

Fershtman, C., and A. Pakes (2012): "Dynamic Games with Asymmetric Information: A Framework for Empirical Work," *Quarterly Journal of Economics*, 127(4), 1611–1661.

Gallego, G., and M. Hu (2007): "Dynamic pricing of perishable assets under competition," in *Agifors reservations and yield management study group annual meeting proceedings.(Jeju Island, South Korea).*

Gallego, G., and G. Van Ryzin (1994): "Optimal dynamic pricing of inventories with stochastic demand over finite horizons," *Management science*, pp. 999–1020.

Gedge, C., J. W. Roberts, and A. Sweeting (2014): "A Model of Dynamic Limit Pricing with an Application to the Airline Industry," Discussion paper, University of Maryland.

Jeziorski, P. (2014): "Empirical Model of Dynamic Merger Enforcement–Choosing Ownership Caps in US Radio," Discussion paper, UC Berkeley.

Klenow, P. J., and B. A. Malin (2010): *Handbook of Monetary Economics*vol. 3, chap. Microeconomic Evidence on Price-Setting, pp. 231–284. Elsevier.

Levin, Y., J. McGill, and M. Nediak (2009): "Dynamic pricing in the presence of strategic consumers and oligopolistic competition," *Management science*, 55(1), 32–46.

Lin, K., and S. Sibdari (2009): "Dynamic price competition with discrete customer choices," *European Journal of Operational Research*, 197(3), 969–980.

Maskin, E., and J. Tirole (1987): "A theory of dynamic oligopoly, III: Cournot competition," *European Economic Review*, 31(4), 947–968.

——— (1988a): "A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs," *Econometrica: Journal of the Econometric Society*, pp. 549–569.

——— (1988b): "A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles," *Econometrica: Journal of the Econometric Society*, pp. 571–599.

McAfee, P., and V. Te Velde (2006): "Dynamic pricing in the airline industry," *forthcoming in Handbook on Economics and Information Systems, Ed: TJ Hendershott, Elsevier.*

Nakamura, E., and J. Steinsson (2013): "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications," *Annual Reviews of Economics*, 5(1), 133–163.

Noel, M. D. (2007): "Edgeworth price cycles, cost-based pricing, and sticky pricing in retail gasoline markets," *The Review of Economics and Statistics*, 89(2), 324–334.

——— (2008): "Edgeworth price cycles and focal prices: Computational dynamic Markov equilibria," *Journal of Economics & Management Strategy*, 17(2), 345–377.

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.

Pang, Z., O. Berman, and M. Hu (2015): "Up Then Down: Bid-Price Trends in Revenue Management," *Production and Operations Management.*

Perakis, G., and A. Sood (2006): "Competitive multi-period pricing for perishable products: A robust optimization approach," *Mathematical Programming*, 107(1-2), 295–335.

Rust, J. (2000): "Parametric Policy Iteration: An Efficient Algorithm for Solving Multidimensional DP Problems?," .

Ryan, S. P. (2012): "The costs of environmental regulation in a concentrated industry," *Econometrica*, 80(3), 1019–1061.

Snider, C. (2009): "Predatory incentives and predation policy: the American Airlines case," Discussion paper, UCLA.

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.

——— (2013): "Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry," *Econometrica*, 81(5), 1763–1803.

Taylor, J. B. (1980): "Aggregate dynamics and staggered contracts," *The Journal of Political Economy*, pp. 1–23.

Wallner, K. (1999): "Sequential moves and tacit collusion: reaction-function cycles in a finite pricing duopoly," *Journal of Economic Theory*, 84(2), 251–267.

Wang, Z. (2009): "(Mixed) strategy in oligopoly pricing: evidence from gasoline price cycles before and under a timing regulation," *Journal of Political Economy*, 117(6), 987–1030.

Zhu, J.-D. (2014): "Effect of Resale on Optimal Ticket Pricing: Evidence from Major League Baseball Tickets," Discussion paper, Texas A&M.