

THE WELFARE EFFECTS OF TICKET RESALE*

Phillip Leslie

NBER, and

Graduate School of Business

Stanford University

pleslie@stanford.edu

Alan Sorensen

NBER, and

Graduate School of Business

Stanford University

asorensen@stanford.edu

November 2007

*We are grateful to Ticketmaster and StubHub for providing data, and to Amitay Alter, Anna Mastri, and Tim Telleen-Lawton for many hours of outstanding research assistance. We are responsible for any errors.

1 Introduction

Active resale markets exist for a wide variety of products: e.g., cars, boats, collectibles, home electronics, sports equipment, musical instruments, textbooks, and music CDs. From a social welfare perspective, economists traditionally regard such markets as unambiguously beneficial, since resale transactions are presumably Pareto-improving exchanges. In general, however, understanding the distributional effects of resale markets is not so straightforward. Resale does not just affect final allocations, it also affects initial allocations—buyers’ choices in the primary market are influenced by the existence and characteristics of the resale market.

Our objective in this paper is to quantify the welfare effects of resale in the market for event tickets, where resale markets are very active and sometimes controversial. Using a rich dataset we have compiled for rock concert tickets, we develop and estimate an equilibrium model in which primary and secondary market outcomes are jointly determined. Resale activity depends on allocations that arise from the primary market, which in turn depend on buyers’ expectations about what will happen in the resale market. Importantly, because our model explicitly incorporates the interplay between the two markets, it allows us to do more than simply say whether (and how much) aggregate welfare is increased by resale markets. Our primary focus is on describing how resale markets reallocate economic surplus among the various market players: sellers, brokers, and consumers.

The policies of primary market sellers suggest there may be some ambiguity about the impact of resale markets on primary market profits. Many sports teams and music artists actively try to prevent resale for tickets to their own events, and have lobbied for strict regulation (if not prohibition) of ticket resale more generally. Others have embraced ticket resale, some going so far as to facilitate resale directly through online ticket exchange services.

Questions also arise regarding the role of brokers, or intermediaries, in resale markets. Some states, like New York, have recently relaxed restrictions on ticket resale while maintaining a requirement that large-volume brokers must be licensed. Conceptually, brokers may be good for social surplus because they add liquidity to the market. However, brokers also extract surplus for themselves, possibly reducing the surplus of the consumers who ultimately attend the events (i.e., the fans). That is to say, brokers may increase the size of the pie, but extract so much themselves that consumers are worse off.

Our dataset covers over 100 rock concerts from the summer of 2004. For each concert, we observe detailed information on primary market ticket sales combined with information on resale transactions. The primary market data come from Ticketmaster, the dominant ticketing agency in the industry, and the secondary market data come from eBay and StubHub, the two leading

online outlets for ticket resale. To our knowledge, this is the first time a dataset has enabled a systematic study of primary and secondary markets in parallel. Even from a purely descriptive standpoint, our analysis of these data constitutes a significant contribution to our understanding of ticket resale.

Concert ticket markets are an attractive setting in which to study the economics of resale because we can abstract from a number of complicating factors that would arise in other resale markets. For example, in the case of ticket resale there is no distinction between new and used versions of the product, as there is for durable goods such as cars. Also, ticket quality is a combination of event quality and seat quality, and since it is plausible that event quality is equally uncertain to both the buyer and seller, and seat quality is given by the seat location which is stipulated on the ticket, we can ignore the important role of adverse selection that arises in most durable goods resale markets.¹

Another advantage of studying event tickets is the relatively simple timing: to a first approximation, primary and secondary markets occur sequentially instead of simultaneously. Markets open and close on specified dates (the ticket onsale date, and the event date, respectively), and while almost all primary market sales occur shortly after the onsale date, most resale activity occurs later (i.e., shortly before the event). This means we can reasonably treat the markets as separated in time, which greatly simplifies the analysis.

Yet another advantage of studying concert tickets is that the seller can plausibly be viewed as a monopolist in the primary market. Hence, we are able to focus on the relationship between primary and resale markets, without incorporating the complexity of rivalry among primary market sellers.

Our structural model incorporates several important aspects of ticket resale. We model two types of buyers, brokers and non-brokers, with the distinction being that brokers get zero utility from consuming a ticket. As mentioned above, buyers' decisions in the primary market are influenced by their expectations about the secondary market, so that primary and secondary market outcomes are jointly determined. Our model requires buyers' expectations to be rational: in equilibrium, secondary market outcomes are on average consistent with the expectations that led to them.

We allow for various frictions in the secondary market, including transaction costs (which may be different for brokers vs. non-brokers) and random participation in resale auctions. We also incorporate various sources of buyer uncertainty. Buyers cannot perfectly forecast secondary market outcomes because they are uncertain about the allocation of tickets in the primary

¹Fraud may still be a concern for ticket buyers in resale markets. However online resale sites such as eBay and StubHub address this concern by offering guarantees of various kinds.

market (i.e., the order in which buyers arrived), about unanticipated schedule conflicts that may arise for non-brokers, and about the overall level of demand. Our estimates allow us to quantify the relative importance of these sources of uncertainty.

The paper proceeds as follows. In section 2 we briefly outline the relevant institutional details about the market for concert tickets. In section 3 we describe how we compiled the data, and provide summary statistics and descriptive analyses. Section 4 outlines the model. Section 5 explains some details about the estimation, and reports some preliminary results. Section 6 [not yet written] discusses the results of various counterfactual simulations, and Section 7 [not yet written] concludes.

2 Industry background

Performances of live music in the U.S. generate roughly \$3 billion in ticket revenue annually, and this number has been steadily increasing for at least the past decade. Concerts are organized and financed by promoters, but the performing artists are principally responsible for setting prices. Typically the artist and/or artist manager consults with the promoter and venue owner to determine the partitioning of the venue and the prices for each partition.

Promoters employ ticketing agencies to handle all the logistics of ticket selling. The dominant firm in this industry is Ticketmaster: it serves as the primary market vendor for over half of the major concerts in the U.S. Ticketmaster sells tickets primarily online or by phone, but also has some physical outlets located in music stores or malls. Tickets usually go on sale three months before the event, and they sometimes sell out very quickly (in some cases within hours of going on sale).

Concert tickets are commonly resold at prices well above face value, which suggests that concerts are systematically underpriced. There are a variety of reasons why artists might choose to underprice their tickets: to generate good will with their fans and avoid appearing greedy; to ensure that the venue is full (which arguably makes for a better concert experience); or to address concerns about fairness. (Artists often cite a desire to give everyone, rich and poor, the same chance at getting a ticket.) In this study, we do not concern ourselves with rationalizing artists' price-setting decisions. Instead, we take the primary market prices as given, and focus our analysis on how resale markets reallocate economic surplus.

Ticket brokers obviously play a big role in resale markets, especially for large events. Brokers have various ways of obtaining tickets. Some allegedly use schemes for bypassing Ticketmaster's

online security measures.² Some brokers employ “pullers” to purchase tickets on their behalf. Others may have relationships with the artist or promoter.³

Professional brokers are not the only ones who resell tickets, however. Many tickets are resold by individuals who purchased tickets with the intention of using them, but subsequently learned they could not attend the event. Importantly, the ability to buy and resell on the internet makes it fairly easy for anyone to act as a broker—e.g., a college student in New York can buy tickets to a concert in Los Angeles through Ticketmaster’s website and then sell those tickets for a profit on eBay, all without leaving her dorm room.

As mentioned above, many states have laws forbidding or regulating certain types of resale. In most cases, however, these anti-scalping laws are loosely enforced, and in any case there are various ways to get around them.⁴ In some ways, these laws seem less like purposeful implementations of policy, and more like expressions of public sentiment toward ticket scalping. Scalpers are held in low regard almost everywhere, but the degree of disdain varies internationally. In some cultures, scalpers are considered among the lowest of life forms; in others (arguably including the U.S.), they are seen as providing a useful (if somewhat distasteful) service.

3 Data

Our data combine detailed information about primary and secondary market sales for a sample of rock concerts performed during the summer of 2004. Our sample is not intended to be representative of the thousands of concerts performed that summer; rather, it focuses on large concerts by major artists, for which resale markets tended to be most active.

For each concert, we observe detailed information about tickets sold through the primary market vendor (Ticketmaster), as well as information about tickets that were resold on eBay or StubHub. To our knowledge, we are the first scholars to analyze primary and secondary markets in parallel. We suspect that data from either market in isolation would still be interesting and informative. But by combining the primary market sales with data on resales, we are able to study the

²Ticketmaster imposes a limit on how many tickets can be purchased in a single transaction, and enforces this by requiring nontrivial input from the buyer as part of the online purchasing process. But some brokers have apparently been able to automate the online purchasing process, allowing them to complete a large number of separate purchases in the time it takes an ordinary buyer to complete a single transaction.

³Artists and promoters typically receive an allocation of tickets, which they can distribute or sell as they wish. Our understanding is that these tickets are often just sold through Ticketmaster along with the other tickets. Anecdotally, however, we were told that occasionally an artist will arrange to sell his allocation of tickets to a broker, who in turn resells the tickets to the public.

⁴See Elfenbein (2005). In our data, we occasionally observe tickets being sold at or under face value, but with exorbitant shipping fees.

interaction between the two markets, which we believe is crucial to answering the relevant economic questions.

Figures 1 and 2 illustrate the data for four example concerts: two by Kenny Chesney and two by the Dave Matthews Band. In each of these figures, the vertical axis represents price, and the horizontal axis represents seat quality (ordered from worst to best). (We explain our measure of seat quality in more detail below.) Consider the first panel of figure 1, which shows the data for a Kenny Chesney concert performed in Tacoma, WA. The dots along the horizontal lines represent tickets that were sold in the primary market, at three different price points. The other dots and squares represent resales by non-brokers and brokers, respectively. Resale activity was concentrated among the highest-quality tickets, and the average premium paid for these tickets was substantial. The figure also illustrates that resale prices were highly variable, with some relatively high-quality seats even being sold below face value.

In the following subsections we explain how the data shown in Figures 1 and 2 were assembled. We describe the primary market and secondary market data sources in turn, and then explain how they were merged. We also report basic summary statistics, and describe some patterns in the data that motivate various aspects of our empirical model.

3.1 Primary market data

The primary market data were provided by Ticketmaster. For each concert, we obtain information from two sources: a “seat map” and a daily sales audit. The seat maps essentially list the available seats at a given event, indicating the order in which the seats were to be offered for sale, and the outcome (i.e., sold, comped, or open).⁵ The daily audits contain ticket prices (including the various Ticketmaster fees), as well as how many tickets were sold in each price level on each day. Essentially, the daily audits allow us to assign prices and dates of sale to the seats listed in the seat maps.

We use the ordering of seats in the seat map data as our measure of relative seat quality. The main virtue of this approach is that it reflects the primary market vendor’s assessment of quality: Ticketmaster uses this ordering to determine the current “best available” seat when a buyer makes an inquiry online or by phone. Also, it allows us to measure quality separately for each seat, as opposed to using a coarser measure (e.g., assigning qualities by section). The seat orderings are also fairly sophisticated: for example, seats in the middle of a row might be preferred over seats toward the outsides of rows further forward, and seats at the front of

⁵An open seat is an available seat that went unsold. The seat maps also identify seats that were “held” or “killed.” These seats were unavailable for sale, so naturally we ignore them in the analysis.

second-level sections are sometimes preferred to seats at the back of first-level sections.

Nevertheless, using this measure of seat quality has its drawbacks. Although the orderings appear to be carefully determined, we suspect they are not always perfect. More importantly, Ticketmaster’s ordinal ranking of tickets does not tell us anything about absolute differences in quality between seats. We know that a given ticket is supposed to be better than all subsequent tickets in the ordering, but we do not know how *much* better. Because we think any information we could bring to bear on absolute quality differences would inevitably be arbitrary, in the analyses below we simply assume that quality differences are uniform—i.e., the difference in quality between seats j and $j + s$ is the same regardless of j . Specifically, we use $1 - (j/J)$ as our measure of quality, where j is the seat’s position in the “best available” order, and J is the total number of tickets available. The best seat ($j = 0$) therefore has quality 1, and the worst seat has quality $1/J$.

Many of the events in our sample have large lawn sections, in which seats are first-come, first-served instead of pre-assigned. Tickets in these sections have identical quality *ex ante*: at the time of purchase, there is no reason that one lawn ticket is any better than another. We assign all lawn tickets the same relative quality for a given event, setting the quality based on the median lawn ticket. That is, if there are 2,000 lawn seats, and 4,000 seats that precede the lawn in the best available ordering, then the lawn seats are all assumed to be at position 5,000 in the ordering.

Complementary tickets (“comps”) represent a potential difficulty for our analysis. Although the Ticketmaster data identify exactly which tickets were given away as comps, we cannot typically tell if a resold ticket was originally a comp, because usually for resales we only know the section and row, not the exact seat number. This is problematic, because we suspect resale is more likely for comps than for ordinary sales.⁶ For some of the concerts in the original sample, a significant fraction of seats was given away, apparently in an effort to fill seats.⁷ For most concerts, however, comps represent a small fraction of total tickets, and tend to be located toward the front—suggesting they are perhaps targeted toward friends of the artist and/or promoter. In order to minimize the potential difficulties associated with comps, we focus on 103 events for which comps account for less than 3% of total tickets. These 103 concerts were performed by 18 distinct artists,⁸ and in total there were 1,739,346 tickets sold or comped.

⁶The reason is that people who purchase tickets value the tickets above the price paid, whereas recipients of comps may have valuations well below the face value. It is possible, however, that comps are successfully targeted at the artist’s strongest fans, in which case they may be *less* likely to be resold.

⁷For example, most of the concerts by Jessica Simpson had several thousand low-quality seats given away as comps.

⁸The artists are Aerosmith, Dave Matthews Band, Eric Clapton, Jimmy Buffett, John Mayer, Josh Groban, Kenny Chesney, Kid Rock, Madonna, Phish, Prince, Rush, Sarah McLachlan, Shania Twain, Sting, Tim McGraw, Usher, and Van Halen.

3.2 Secondary market data

To obtain information about resales, we captured and parsed eBay auction pages for all tickets to major concerts in the summer of 2004. From these pages we determined how many tickets were sold, on what date, at what price (including shipping), and the rough location of the seats. We only use auctions that ended with a sale (either via a bid that exceeded the seller’s reserve, or via “Buy-it-now”). The auction pages also list information about the seller, including the seller’s eBay username. We use this to distinguish between brokers and non-brokers: we categorize an eBay seller as a broker if we observe him selling 10 or more tickets in the data.

We also obtained data from StubHub, a leading online marketplace designed specifically for ticket resale. For every concert in our sample, we observe all tickets sold through StubHub, and for each transaction we observe the number of tickets sold, the seat location, the price (including shipping and fees), the date, and the seller identity and classification (broker vs. non-broker).

Matching eBay auctions to specific concert events was straightforward, albeit tedious, because the auction pages contain standardized fields for the venue and event date. The process of assigning resales to specific seats was more complicated, because exact seat numbers were rarely reported in the eBay or StubHub auctions. We were able to determine the section and row for 75% of the resale transactions; for another 23% we could only determine the section. For the remaining 2%, our parser did not even detect the section, and we simply dropped these transactions from the analysis. We are left with 68,828 resold tickets (the vast majority of them on eBay).

Beginning with transactions for which we observed both the section and row, we assigned resales to specific seats by spreading them evenly throughout the relevant section and row. So, for example, if in the Ticketmaster data we see that there were 20 seats in section 101, row 3, and we observe 3 tickets resold on eBay or StubHub in section 101, row 3, we assign them to seats 5, 10, and 15 within that row. Once the section-and-row transactions were assigned, we then assigned section-only transactions using the same principle. Suppose that after assigning all the section-and-row transactions, 200 seats in section 101 remain unassigned. Then if we observe 4 tickets resold in section 101, unknown row, we assign those 4 tickets to seats 40, 80, 120, and 160 (of the 200).⁹

For the empirical model we estimate below, it would be ideal to observe *all* resale activity for the sample concerts. Many tickets are presumably resold through the various broker websites

⁹We also tried putting resales at the middle of their respective sections and rows, instead of spreading them evenly. That is, if there were 20 seats in the row, and three tickets resold, we assign the resales to seats 9,10, and 11 instead of 5, 10, and 15. We will check to make sure our empirical results are robust to this alternative approach.

(e.g., razorgator.com, buyselltix.com) or by non-brokers using classified ads in a local newspaper or craigslist.com. Unfortunately, we do not know exactly how much of total resale activity is accounted for by eBay and StubHub. Nor can we directly verify that eBay and StubHub are representative of resale activity more broadly. However, eBay is widely assumed by people in the industry to be the largest single outlet for ticket resales, with StubHub likely the second largest. This was verified by Alan Krueger in a survey of concertgoers at a major rock concert in 2005; he found that eBay and StubHub accounted for between a third and a half of all resold tickets. The fact that both brokers and non-brokers have a significant presence on eBay suggests that our data might be roughly representative of resale activity on broker sites or on craigslist.com.

3.3 Summary statistics

As mentioned above, the dataset covers all 1.7 million tickets sold in the primary market for 103 concerts by 18 different artists. Table 1 provides more detailed summary statistics. The capacity of each concert varies from 3,171 to 35,062 (the median is 15,970). The events tended to sell-out: 47% of concerts sold 100% of their capacity, and 76% sold over 98% of their capacity.

The average ticket price in the primary market was \$83.77.¹⁰ However, there is a good deal of price variation across events: the inter-quartile range of the distribution of average prices across events is from \$54.48 to \$91.74. There is also within-event price variation. Most events tend to offer tickets at about 4 different price levels (varying with seat quality). The maximum number of price levels for a single event in our data is 12, and 2 concerts sold all tickets at a single price. Figure 2 depicts one of these events with a single price in the primary market (in this case, for all 24,873 seats). As we show in our analysis, the low number of price levels in the primary market, relative to capacity, is a key driver of resale activity. The highest price observed in the primary market for any event is \$316.15.

We observe about 69,000 of the tickets purchased in the primary market being resold at eBay or StubHub (i.e. 3.9% of the number of tickets). As shown in Table 1, the maximum number of tickets resold for a given event is 3,130, or 17% of the tickets sold in the primary market. For most events the fraction of tickets resold is between 2% and 5%. On average, total revenue to resellers is equivalent to 6% of the primary market revenue, and the maximum for any single event is a striking 37%. It is important to remember that these numbers are based on reselling at eBay or StubHub alone, which provides a lower bound for the total amount of resale activity.

Table 2 provides additional summary statistics of resale activity.¹¹ The average resale price is

¹⁰Includes shipping and any fees. To be clear, \$83.77 is the average across events of the average ticket price at each event. If all tickets are weighed equally, the average ticket price in the primary market is \$77.20.

¹¹In Table 1 an observation is an event, while in Table 2 an observation is a resold ticket.

\$111.66. Resellers are making significant profits from this activity: the average markup is 40% over the face value, and 25% of resold tickets obtain markups above 66%. On the downside for resellers, 28% of tickets are sold below face value. Resold tickets are not a random sample of those purchased in the primary market, and in particular the resold tickets tend to have significantly higher face values than non-resold tickets: the average price in the primary market of the resold tickets is \$89.71, while the average price of all purchased tickets in the primary market is \$77.20.

One reason why resold tickets tend to have relatively high face values is that resold tickets tend to be for relatively better seats. In Table 2 we report the average seat quality of resold tickets is 0.62 (median is 0.69). In comparison, the average seat quality of tickets purchased in the primary market is 0.51. However, a sizable fraction of resold tickets are also for relatively low quality seats: 25% of resold tickets have seat quality below 0.38. Figure 3 provides a more complete picture of the tendency for resold tickets to be higher quality than resold tickets. The figure shows the predicted probability of resale as a function of seat quality, obtained either from a parametric regression of a resale indicator (equal to one if the ticket was resold) on a cubic polynomial in the seat quality variable, or from a semiparametric regression.¹² It is clear that the higher the seat quality the more likely the ticket will be resold.

Of course seat quality is a key determinant of prices in both the primary and secondary markets. But there are a couple of important differences between these markets in terms of how price relates to seat quality. First, in the primary market prices are based on coarse partitions of each venue (as discussed above). Meanwhile, resale prices reflect small differences in seat quality—every seat may have a different price. Figure 4 depicts the general pattern of resale prices as a function of seat quality (showing both parametric and non-parametric functions). It is apparent that resale prices vary significantly according to seat quality. This is especially true for about the 20% of highest quality seats, where resale prices are a particularly steep function of seat quality. This figure also serves as a reality check on the data.

A second important difference in how seat quality is priced in the primary and secondary markets relates to monotonicity. Primary market prices are weakly monotonically increasing in seat quality for a given event. In contrast, Figures 1 and 2 illustrate that resale prices are a rather noisy function of seat quality, and there are numerous instances of a low quality seat resold at a higher price than a higher quality seat (for a given event). This is basic evidence of inefficiency in the resale market. On the one hand, the resale market allows price to be a more flexible function of seat quality. On the other hand, some form of friction in the resale market causes significant variance in price conditional on seat quality. As we detail in the next

¹²In both cases, event fixed effects are included. The semiparametric regression is an adaptation of Yatchew's (1997) difference-based estimator for partial linear regression models.

section, our empirical model explains this fact as being a consequence of limited participation by potential buyers in resale market auctions.

The analysis in this study emphasizes the consequences of limited price flexibility in the primary market (i.e. utilizing only a few price levels) on resale activity. In Figure 5 we present basic evidence in support of this view. By definition, all seats in a given price level at a given event have the same face value. However, there can be thousands of seats in a given price level, and the difference in seat quality between the best and worst seats in the price level can be dramatic. At equal prices there will be higher demand for the good seats in a given price level than the bad seats. We therefore expect more resale activity for the relatively good seats in any given price level. Figure 5 shows exactly this pattern.

Figure 5 is evidence that unpriced seat quality is an important driver of resale activity. But notice also that the lowest quality seats in any given price level are also resold with positive probability (roughly 2%). This is consistent with other drivers of resale activity, such as general underpricing or schedule conflicts. From the figure it appears that unpriced seat quality may be the most important driver of resale activity. In the analysis below we quantify the relative importance of these different factors.

In the primary market tickets typically go onsale 3 months before the event date. In Table 1 we report that (averaged across events) 64% of tickets purchased in the primary market occur in the first week. In fact the concentration of sales in the first week is even more striking than this number suggests. In the top panel of Figure 6 we depict the complete time-pattern of sales in the primary market. It is clear that sales in the primary market are highly concentrated at the very beginning. The time-pattern of sales in the resale market is less concentrated than the primary market, as shown in the lower panel of the figure. In Table 2 we report that 50% of resale transactions occur within 24 days of the event, and in fact 26% of resale transactions are within 7 days of the event. In the model presented in the next section we assume primary market transactions occur in period 1, and resale transactions occur in period 2. The above facts suggest this is a reasonable simplification.

As noted above, brokers are potentially important agents in resale markets. Based on seller identifiers, 11% of the sellers in the resale market are brokers, and they account for 55% of resold tickets (as reported in Table 2). In the table we also report that 23% of the tickets resold by brokers were at prices below the purchase price in the primary market. By comparison, 33% of the tickets resold by non-brokers were sold below face value. One possible explanation for this difference is that unlike brokers (who never intend to attend the event), non-brokers sometimes resell tickets because of schedule conflicts. There are other possible explanations, but these numbers suggest 10% of the resale activity by non-brokers is an upper bound for the fraction

of resales that are due to schedule conflicts.¹³ In the structural analysis we obtain a specific estimate of the rate of schedule conflicts (taking into account other specified reasons for resale).

Lastly, the total profit (i.e. aggregate markup) obtained from ticket resale in our data is slightly over \$1.5 million. This is equivalent to 1.14% of total revenue in the primary market for these events. As a measure of “money left on the table” this suggests a fairly low degree of forgone profit by firms in the primary market. But this is misleading because modified pricing policies that capture some of this value will also capture more surplus from buyers in the primary market that do not resell tickets. We address this issue in our counterfactuals based on the estimated structural model.

4 Model

We consider a two-period model describing equilibrium behavior in the primary and secondary markets. We assume that all primary market transactions take place exclusively in the first period, and resales take place exclusively in the second period—i.e., there is no overlap in the timing of the two markets. This greatly simplifies the analysis, and appears to be a reasonable approximation of reality given the timing of sales described in Figure 6.

In the present version of the model, we also assume that primary market prices are exogenously given. In a full equilibrium model, the primary market seller’s pricing problem would depend on characteristics of the secondary market. In future research we plan to extend the model in this direction.

4.1 Basic setup

There are two types of agents in our model: brokers and consumers. The principal distinction between them is that brokers get no utility from consuming a ticket: if they purchase in the primary market, it is only with the intention of reselling at a profit. We assume there are M potential buyers in the market, a fraction β of which are brokers.

In the first period, consumers and brokers make their primary market purchase decisions. To capture the stochastic nature of ticket availability in the primary market (e.g., due to thousands of people calling or visiting the Ticketmaster website at exactly 10:00AM when the tickets go on sale), we assume that buyers make their decisions in a random sequence. Let $g_z(z)$ be the

¹³Another possible explanation is that brokers are better at identifying events where they are less likely to resell tickets at a loss.

probability of a given sequence z . In the present version of the model, we assume this probability is independent of buyer characteristics; later we will explore the consequences of letting arrival be correlated with willingness to pay and/or whether the buyer is a broker. Buyers are limited to choosing from the set of unsold tickets at their turn in the sequence, and each buyer is limited to buying one ticket.¹⁴

Consumers vary in their willingness to pay for seat quality. Let $\nu_j \in (0, 1]$ denote the quality of ticket j , measured as described in section 3.1.¹⁵ Consumer i 's gross utility from attending event k in seat j is

$$v^c(\omega_i, \nu_j) \equiv \mu_k(1 + \omega_i \nu_j^\phi). \quad (1)$$

The ω_i term represents the consumer's willingness to pay for quality. The chosen functional form gives us a natural way to interpret ω : the ratio of a consumer's willingness to pay for the best seat ($\nu_j = 1$) vs. the worst seat ($\nu_j = 0$) is just $1 + \omega$. We assume the ω_i 's are drawn from an exponential distribution with mean λ . The curvature term, *phi*, captures the potential nonlinearity of premia for high quality seats (as evidenced in Figure 4). The idea is that even for a given consumer, willingness to pay is likely to be a nonlinear function of seat quality. Event-specific variation in willingness to pay is captured by μ_k .

In the second period, brokers and consumers holding tickets may sell those tickets via auction, where the buyers in the auction are consumers who chose not to purchase (or were rationed) in the primary market. Both brokers and consumers incur transaction costs if they choose to sell. We denote these costs τ^b and τ^c , respectively. Of course, consumers also have the option of using the ticket they purchased in the primary market. Thus, in the resale market the sellers' reservation prices depend on their transaction costs and (for consumers) on their consumption utilities.

4.2 Clearing the resale market

A natural way to clear the resale market would be to calculate every buyer's willingness to pay for every ticket (with the ticketholder's willingness to pay being equal to her reservation price), and then find a vector of prices such that there is no excess demand for any ticket. This is the kind of approach that is typically used to find equilibria in two-sided one-to-one matching

¹⁴We do not model the purchaser's choice of how many tickets to buy. It might be more realistic to assume that the unit of purchase is a pair of tickets, but we doubt this would meaningfully affect our results and conclusions.

¹⁵If j is the ticket's position in the "best available" order, and there are a total of J available, then $\nu_j \equiv 1 - (j/J)$.

markets, for example. Although this approach is feasible in our model, it predicts resale prices that are monotonic in seat quality, which is clearly not true in the data. Observed resale prices increase on average as a function of seat quality, but there is considerable variance in prices conditional on seat quality.

To accommodate this feature of the data, we clear the resale market using a sequence of auctions with incomplete bidder participation. We begin with the highest quality ticket. From the pool of potential buyers (i.e., consumers who did not obtain a ticket in the primary market, and who did not have a schedule conflict), we randomly select L bidders. The owner of the ticket is offered a price equal to the second-highest willingness to pay among those L bidders. If the offer exceeds the owner's reservation price, then the ticket is transacted at that price: the bidder with the highest willingness to pay gets the ticket, and both seller and buyer exit the market.¹⁶ If the offer is below the reservation price, the ticket remains with the seller. (If the seller is a consumer, she uses the ticket herself and gets the utility defined in equation 1. If the seller is a broker, she gets utility zero.) Losing bidders remain in the pool of potential buyers. This process is then repeated for all tickets that were purchased in the primary market, in order of decreasing quality.

4.3 Equilibrium

Buyers' decisions in the primary market clearly depend on their expectations about the resale market. There are four sources of uncertainty about outcomes in the resale market. The first is randomness in the arrival sequence, z , as mentioned above. Unless the resale market is entirely frictionless ($\tau^b = \tau^c = 0$), the equilibrium will depend on the allocation of tickets in the primary market, which in turn depends on the order in which buyers made their purchase decisions. A second source of uncertainty that we introduce is the possibility of unanticipated schedule conflicts. Formally, we assume there is a probability ψ that a given consumer will have zero utility from attending the event, with the uncertainty being resolved in between periods 1 and 2. Notice that if ψ is large, the ability to resell tickets in a secondary market may significantly increase willingness to pay in the primary market. Also, note that resale market outcomes depend on *which* consumers have schedule conflicts. For example, if by chance the individuals with conflicts are predominantly ones with high willingness to pay, then their absence from the resale market will tend to make prices lower. Because the composition of schedule conflicts matters, we denote Ψ to be an $M \times 1$ vector of indicators identifying which consumers have conflicts.

¹⁶We allow only one transaction per period for any individual. So we do not allow consumers to buy in the primary market, sell in the resale market, and then buy another ticket in the resale market.

Randomness in auction participation is the third source of uncertainty. As explained above, we clear the secondary market using a sequence of auctions, with only a random subset of potential buyers participating in each auction. Obviously, realized outcomes in the resale market will depend on the particular subsets of buyers who bid for each ticket. We define H to be an $MJ \times 1$ vector containing a random ordering of the M buyers for each of the J tickets. This can be thought of as the order in which buyers “arrive” at each of the secondary market auctions, with only the first L arrivals being allowed to participate.

The fourth source of uncertainty is about the distribution of ω , consumers’ willingness to pay for quality. Buyers know their own ω ’s (if they are consumers), and they know the distribution of ω ’s is exponential, but we assume they do not know λ —i.e., buyers are uncertain about the mean of the distribution of willingness to pay. Buyers believe that λ is drawn from some distribution with density g_λ . To keep things simple, we assume consumers do not use their own ω ’s as signals with which to update their beliefs about λ .

Incorporating this kind of uncertainty is necessary if we want the model to fit the data. Specifically, for many events we observe both consumers and brokers reselling tickets below face value. For consumers, such transactions could be explained by unanticipated schedule conflicts; but for brokers, we would never observe resales below face value unless brokers sometimes overestimate the strength of demand. Essentially, this source of underlying uncertainty allows us to explain why some events sell out in the primary market but then have very thin resale markets with very low prices, while other events do not sell out in the primary market but then have very high prices in the resale market.

The price of a ticket in the resale market is principally a function of its quality, but it will also depend on the realizations of the random variables described above. For notational convenience, we define the resale price function $R(\nu|z, \Psi, H, \lambda)$, which we will sometimes write simply as $R(\nu|\cdot)$.

The decision problem for a broker in the first period is straightforward. Her strategy is to purchase the ticket j that maximizes

$$E(u_j^b) = E(R(\nu_j|\cdot)) - p_j - \tau^b ,$$

where p_j is the primary market price of ticket j , and $E(R(\nu_j|\cdot))$ is the expected price of ticket j in the resale market, where the expectation is with respect to the four sources of uncertainty described above. Of course, if the transaction cost τ^b exceeds the expected resale profits, a broker also has the option of not purchasing a ticket.

A consumer's decision problem is somewhat more complicated, as illustrated in Figure 7. If a consumer buys ticket j in the primary market, with probability ψ she will be forced to resell the ticket, obtaining some price $R(\nu_j|\cdot)$. (While not illustrated explicitly in the figure, she also has the option of discarding the ticket if the transaction cost is higher than the resale profit, in which case her payoff is $-p_j$.) If she has no schedule conflict, she will have the choice of reselling or using the ticket, with the latter option delivering a net utility of $\mu_k(1 + \omega_i\nu_j^\phi) - p_j$. The expected payoff from buying ticket j is therefore

$$E(u^c|\text{buy } j) = -p_j + \psi E(\max\{0, R(\nu_j|\cdot) - \tau^c\}) + (1 - \psi) E(\max\{0, R(\nu_j|\cdot) - \tau^c, \mu_k(1 + \omega_i\nu_j^\phi)\}) . \quad (2)$$

where again the expectations are with respect to z , Ψ , H , and λ .

If instead the consumer chooses not to buy a ticket in the primary market, but rather wait until the secondary market, her expected utility is given by

$$E(u^c|\text{wait}) = (1 - \psi) E(\max\{0, \mu_k(1 + \omega_i\tilde{\nu}^\phi) - R(\tilde{\nu}|\cdot)\}) . \quad (3)$$

Here, the consumer is not only uncertain about what prices will be in the resale market, she is also uncertain about which ticket (if any) she will be able to buy in the resale market. We use the notation $\tilde{\nu}$ to indicate that ticket quality is itself a random variable for a consumer who chooses to delay her purchase until period 2.

Given this payoff structure, a rational expectations equilibrium is one in which (i) brokers and consumers make their first-period decisions optimally given their expectations about second-period outcomes; and (ii) those expectations are on average correct given optimal decision-making in the first period. Given a primary market allocation and realizations of the model's random variables (z , Ψ , H , and λ), the resale market outcomes follow deterministically. The challenge is finding expectations that rationalize a set of primary market decisions that in turn lead to resale market outcomes consistent (on average) with those expectations. In other words, the trick is to find a fixed point in the mapping of expectations into average resale market outcomes.

The expectations described in the equations above cannot be calculated analytically, even for particular assumptions about the probability distributions of z , Ψ , H , and λ . Although realizations of these random variables lead deterministically to a set of resale market outcomes, the

form of the function $R(\nu|z, Psi, H, \lambda)$ is not known. Nor is it possible to determine the value of \tilde{v} as a function of z , Ψ , H , and λ .

We therefore take a computational approach to solving this problem. We conjecture a parameterized approximation to the buyers' expected values, and then iterate on the parameters of that approximation until we converge to a fixed point. As explained above, a buyer's expected utility, as a function of the primary market choice, depends on (i) whether the buyer is a broker or consumer; (ii) the quality (ν) of the ticket purchased, if any; and (iii) the buyer's ω if the buyer is a consumer. We therefore choose a parametric function $V(b, \nu, \omega|\alpha)$ to represent buyers' expectations, where b is an indicator for whether the buyer is a broker, and α are the parameters. Our algorithm for finding a fixed point is as follows:

- (1) Choose an initial set of parameters, α_0 . Simulate primary and secondary market outcomes for S draws on the model's random variables (arrival sequences, schedule conflicts, etc.), where consumers make primary market choices to maximize $V(b, \nu, \omega|\alpha_0)$.
- (2) Use the *realized* final utilities from the simulations in step (1) to re-estimate the function $V(b, \nu, \omega|\alpha)$. Essentially, we regress realized utilities on a function of b , ν , and ω to obtain a new set of parameters, α_1 .
- (3) Use the new set of parameters from (2) to simulate primary and secondary market outcomes as in (1). Iterate on steps (1) and (2) until V converges—i.e., until $V(b, \nu, \omega|\alpha_t)$ is sufficiently close to $V(b, \nu, \omega|\alpha_{t-1})$.

In estimating the model below, we use a very simple parameterization of V . Letting h be an indicator for whether the buyer holds a ticket going into the second period, we let

$$V(b, \nu, \omega|\alpha) = b \cdot h \cdot (\alpha_0 + \alpha_1 \nu) + (1-b) \cdot h \cdot (\alpha_2 + \alpha_3 \nu + \alpha_4 \omega + \alpha_5 \nu \omega) + (1-b) \cdot (1-h) \cdot (\alpha_6 + \alpha_7 \omega) . \quad (4)$$

This parameterization captures the essential elements of the expectations described above. For a broker, expected utility depends only on the quality of the ticket owned, ν . For a consumer without a ticket, expected utility depends only on the consumer's willingness to pay for quality, ω . For a consumer holding a ticket, expected utility depends on both ν and ω , since ultimately the ticket will either be consumed (yielding a payoff that depends on ν and ω) or resold (yielding a payoff that depends on ν).

Convergence of this algorithm means we have found a rational expectations equilibrium: a set of expectations V such that the primary market choices that follow from V lead to secondary

market outcomes consistent with V . The convergence criterion we use is based on average differences in V . At each iteration of the algorithm, we essentially estimate the regression described in (4) using $M \times S$ “observations.” We stop iterating when

$$\frac{1}{MS} \sum_{i=1}^{MS} \left(\frac{|V_i(\alpha_t) - V_i(\alpha_{t-1})|}{V_i(\alpha_{t-1})} \right) \leq 0.005 .$$

In other words, we stop when the fitted values of V differ from those of the previous iteration by less than half of one percent on average.

5 Estimation and results

Given a set of event characteristics (e.g., prices and capacities by price level), the model described above allows us to predict both primary market sales and secondary market sales (including resale prices) as a function of the parameters. Heuristically, our estimation approach is simply to find a set of parameters that minimizes the differences between the outcomes we observe in the data and those predicted by the model.

Letting j index the seats in our data, we define four outcome variables:

$$y_{1j} = \begin{cases} 1 & \text{if ticket } j \text{ was purchased in period 1} \\ 0 & \text{otherwise} \end{cases}$$

$$y_{2j} = \begin{cases} 1 & \text{if ticket } j \text{ was resold by a broker in period 2} \\ 0 & \text{otherwise} \end{cases}$$

$$y_{3j} = \begin{cases} 1 & \text{if ticket } j \text{ was resold by a consumer in period 2} \\ 0 & \text{otherwise} \end{cases}$$

$$y_{4j} = \text{resale price of ticket } j, \text{ conditional on resale}$$

Our model allows us to calculate expected values of these dependent variables given the data and parameters—i.e.,

$$E[y_{lj}|x; \theta] = g_l(x; \theta) \text{ for } l = 1 \dots 4$$

We estimate θ using GMM, with the above serving as conditional moment restrictions. We define

$$m_l = \frac{1}{N} \sum_{j=1}^N x_j (y_{lj} - g_l(x_j; \theta))$$

where x_j is a $K \times 1$ vector of instruments, and N is the number of tickets in the entire dataset. Letting $m \equiv [m'_1 m'_2 m'_3 m'_4]'$ be the stacked vector of moment restrictions, we obtain an estimate of θ by minimizing $m'Am$, with A being the variance-minimizing weight matrix suggested by Hansen. The instruments x include a constant, event dummies, and seat quality. Since our data contain 103 events, we have $104 \times 4 = 416$ moment restrictions, so the parameters of the model are overidentified.

5.1 Estimation details

The main computational burden in estimating the model comes from simulating the primary and secondary market outcomes for a given set of parameters. This simulation needs to be done separately for each event, since events differ in their pricing structures, and, as described in section 4.3 above, in each case we must iterate until we converge to an equilibrium set of beliefs about secondary market outcomes.

In order to make this burden more manageable, instead of simulating outcomes for events with thousands of seats, we simulate events with 200 seats, and then scale up the predictions to match the size of the event in question. For example, for an event with 10,000 seats, with 4,000 and 6,000 seats in two respective price levels, we simulate primary and secondary market outcomes for an event with 200 seats, with 80 and 120 seats in the two respective price levels. We then “scale up” by applying the predictions for seat 1 in the simulated event to seats 1-50 in the actual event, the predictions for seat 2 to seats 51-100, and so on.¹⁷

5.2 Identification

Two important variables in our model are neither known to us as data nor identified by the data as parameters. The first is the size of the market, M . In the estimates reported below, we fix M to be three times the capacity of the event; later we can check whether our results

¹⁷This obviously introduces additional noise into our estimator, but in principle we can eliminate as much of this noise as we want by increasing the size of the simulated event up to the size of the actual event.

are robust to alternative assumptions. The second is the fraction of total resales that our data account for. As mentioned above, we do not have a precise estimate of this number, but the scant evidence we do have suggests it is between 30 and 50 percent. For the results we report below, we assume that eBay and StubHub account for 50 percent of total resales. This factors into the estimation when we match predicted resale probabilities to observed resale outcomes; we simply divide in half the probabilities predicted by the model—i.e., we match the data to the probability of resale times the probability of observing that resale.

In this version of the paper we are also fixing the number of bidders in the resale auctions, L , to be 10. This is something we intend to relax in the future, since in principle there is variation in the data to identify L —namely, the variance in resale prices conditional on seat quality.

For the parameters we do estimate, we can offer a heuristic explanation of how the patterns in the data identify them. The event dummies μ_k are basically identified by differences across events in the overall level of resale prices. The “curvature” parameter ϕ is identified by the shape of the relationship between resale prices and seat quality (e.g., as shown in Figure 4).

The shape of the price-quality relationship also influences $\bar{\lambda}$, the mean of the distribution of ω ’s. However, this parameter is driven primarily by the level of resale prices for the highest-quality tickets: as explained above, a consumer’s ω determines the ratio of her willingness to pay for the best seat vs. the worst seat. If in the data we observe that resale prices for the best seats are typically 3 times more than for the worst seats, then $\bar{\lambda}$ needs to be such that the highest draws of ω are around 2.

The standard deviation of beliefs, σ_λ , is identified by the frequency with which tickets are resold at a loss. Essentially, the more often we observe instances where buyers (especially brokers) overestimated demand for an event, the larger will be our estimate of σ_λ .

The fraction of buyers who are brokers (β) is basically driven by the relative frequency of sales by brokers in the resale market. To be clear, however, the estimate will not simply equal the frequency of brokers in the data. If consumers have higher transaction costs than brokers, as we expect, then brokers will be more likely than consumers to speculate in the primary market—so even a small β could be consistent with a large fraction of resales being done by brokers.

Identification of the transaction costs is driven by ticketholders’ relative propensity to resell at high vs. low expected markups. Loosely speaking, positive transaction costs allow the model to rationalize low rates of resale in the data even for tickets that would have fetched very high markups. Not so loosely speaking, the transaction costs estimates should depend on the slope of the relationship between the probability of resale and the expected markup, and specifically on where that slope becomes positive. For example, suppose that τ^c is equal to \$10. For tickets

that would resell for less than \$10 above face value, the model will predict very low probabilities of resale by consumers. More importantly, the probability of resale will be independent of the expected markup if that markup is less than \$10. Only as the expected markup rises above \$10 will the probability of resale increase—i.e., at \$10 the slope would become positive.

Finally, the probability of schedule conflicts, ψ , is driven by the relative rate at which consumers vs. brokers resell below face value. The model assumes that both types of buyer have the same information, so they should be equally likely to overestimate demand for an event. To the extent that consumers are more likely than brokers to sell at a loss, in the model this must be driven by schedule conflicts (which matter for consumers but are irrelevant for brokers).

5.3 Results

In Table 3, we report estimates of the model for a “starter dataset” containing only a subset of 10 events taken from the larger sample. (The subset includes all four of the events shown in Figures 1 and 2.)

6 Welfare implications

Counterfactual simulations: How is the economic surplus obtained by primary market sellers, brokers, and non-brokers affected if we

- shut down the resale market? (i.e., set $\tau^b = \tau^c = \infty$)
- reduce transaction costs?
- eliminate transaction costs?
- eliminate brokers? (set $\beta = 0$)
- increase participation in the resale auctions? (increase L)
- make primary market pricing more sophisticated?

7 Conclusions

None yet.

References

- Busch, L.A. and P. Curry (2005): “Rock Concert Pricing and Anti-Scalping Laws: Selling to an Input,” *Mimeo*, University of Waterloo and Simon Fraser University.
- Chevalier, J.A. and A. Goolsbee (2005): “Are Durable Goods Consumers Forward Looking? Evidence from College Textbooks,” *Mimeo*.
- Chu, S., P. Leslie and A. Sorensen (2006): “Incomplex Alternatives to Mixed Bundling,” *Mimeo*, Stanford University.
- Connolly, M. and A.B. Krueger (2005): “Rockonomics: The Economics of Popular Music,” *Mimeo*, Princeton University.
- Courty, P. (2003): “Some Economics of Ticket Resale,” *Journal of Economic Perspectives*, 17(2), 85–97.
- Depken, C.A. (2006): “Another Look at Anti-Scalping Laws: Theory and Evidence,” *Mimeo*, University of Texas at Arlington.
- DeSerpa, A.C. (1994): “To Err Is Rational: A Theory of Excess Demand for Tickets,” *Managerial and Decision Economics*, 15, 511-18.
- DeSerpa, A.C. and R.L. Faith (1996): “Bru-u-u-uce: The Simple Economics of Mob Goods,” *Public Choice*, 89, 77–91.
- Elfenbein, D. (2005): “Do Anti-Scalping Laws Make a Difference Online? Evidence from Internet Sales of NFL Tickets,” *Mimeo*, Washington University in St. Louis.
- Jennings, M. and S. Happel (2002): “Creating a Futures Market for Major Event Tickets: Problems and Prospects,” *The Cato Journal*, Winter.
- Jennings, M. and S. Happel (1995): “The Folly of Antiscalping Laws,” *The Cato Journal*, Spring.
- Jennings, M. and S. Happel (1989): “Assessing the Economic Rationale and Legal Remedies for Ticket Scalping,” *Journal of Legislation*, Spring.
- Karp, L. and J. Perloff (2003): “When Promoters Like Scalpers,” *Mimeo*, University of California, Berkeley.
- Leslie, P. (2004): “Price Discrimination in Broadway Theatre,” *RAND Journal of Economics*, 35(3), 520–41.
- Mortimer, J. and S. Sorensen (2005): “Supply Responses to Digital Distribution: Recorded Music and Live Performances,” *Mimeo*, Harvard University and Stanford University.
- Swafford, J.L. (1999): “Arbitrage, Speculation, and Public Policy Toward Ticket Scalping,” *Public Finance Review*, 27(5), 531–40.

Table 1: Summary statistics: Events ($N = 103$)

	Mean	Std. Dev.	Percentiles				
			Min	.25	.50	.75	Max
Primary Market:							
Tickets sold	16679.04	6475.86	3169.00	13148.00	16012.00	19490.00	34844.00
Tickets comped	207.82	150.43	0.00	74.00	192.00	329.00	731.00
Revenue (000)	1287.60	563.59	230.76	892.10	1182.34	1662.83	2668.52
Venue capacity	17329.29	6638.35	3171.00	13800.00	15970.50	19870.00	35062.00
Capacity util.	0.98	0.15	0.39	0.98	1.00	1.00	1.56
Average price	83.77	40.41	39.91	54.48	69.05	91.74	188.64
Maximum price	130.35	100.63	43.50	65.70	85.25	139.00	316.15
# price levels	4.99	2.27	1.00	3.00	4.00	7.00	12.00
% first week	0.64	0.18	0.01	0.52	0.68	0.78	0.96
Secondary Market:							
Tickets resold	668.23	535.64	38.00	334.00	509.00	846.00	3130.00
Resale revenue	74.62	57.60	3.07	32.72	57.97	101.53	295.32
Percent resold	0.04	0.03	0.00	0.02	0.03	0.05	0.17
Percent revenue	0.06	0.05	0.00	0.03	0.05	0.07	0.37

Revenue numbers are in thousands of U.S. dollars. “# price levels” is the number of distinct price points for the event. “% first week” is the percentage of primary market sales that occurred within one week of the public onsale date. “Percent resold” is the number of resales observed in our data divided by the number of primary market sales, and “Percent revenue” is the resale revenue divided by primary market revenue.

Table 2: Summary statistics: Resold tickets ($N = 68,828$)

	Mean	Std. Dev.	Percentiles				
			Min	.25	.50	.75	Max
Resale price	111.66	78.40	3.03	65.00	91.25	134.98	2000.00
Markup	21.95	66.69	-308.65	-3.21	18.97	44.45	1688.85
% Markup	0.40	0.76	-0.98	-0.04	0.29	0.66	13.04
Seat quality	0.62	0.27	0.00	0.38	0.69	0.87	1.00
Days to event	41.89	41.75	0.00	7.00	24.00	73.00	208.00
Sold by broker	0.55	0.50	0.00	0.00	1.00	1.00	1.00
Sold below face value:							
by broker	0.23	0.42	0.00	0.00	0.00	0.00	1.00
by non-broker	0.33	0.47	0.00	0.00	0.00	1.00	1.00

Resale prices include shipping fees. Markups are calculated relative to the ticket’s face value, including shipping and facility fees. Seat quality is based on the “best available” ordering in which Ticketmaster sold the tickets, as explained in the text, and is normalized to be on a $[0,1]$ scale (1 being the best seat in the house). Brokers are eBay sellers who sold 10 or more tickets in our sample, or StubHub sellers who were explicitly classified as brokers.

Table 3: Results

Parameter	Notation	Estimate	Standard Error
Consumers' transaction cost	τ^c	44.358	
Brokers' transaction cost	τ^b	12.954	
Curvature	ϕ	1.631	
Mean of $\log(\lambda)$	$\bar{\lambda}$	-0.035	
SD of $\log(\lambda)$	σ_λ	0.250	
Prob(conflict)	ψ	0.015	
Prob(broker)	β	0.055	
Event fixed effects:			
Event 1	μ_1	55.407	
Event 2	μ_2	51.426	
Event 3	μ_3	49.847	
Event 4	μ_4	65.422	
Event 5	μ_5	59.198	
Event 6	μ_6	40.977	
Event 7	μ_7	42.595	
Event 8	μ_8	46.415	
Event 9	μ_9	65.054	
Event 10	μ_{10}	46.594	

These estimates are from a subsample of 10 events, as explained in the text.

Figure 1: Two sample events

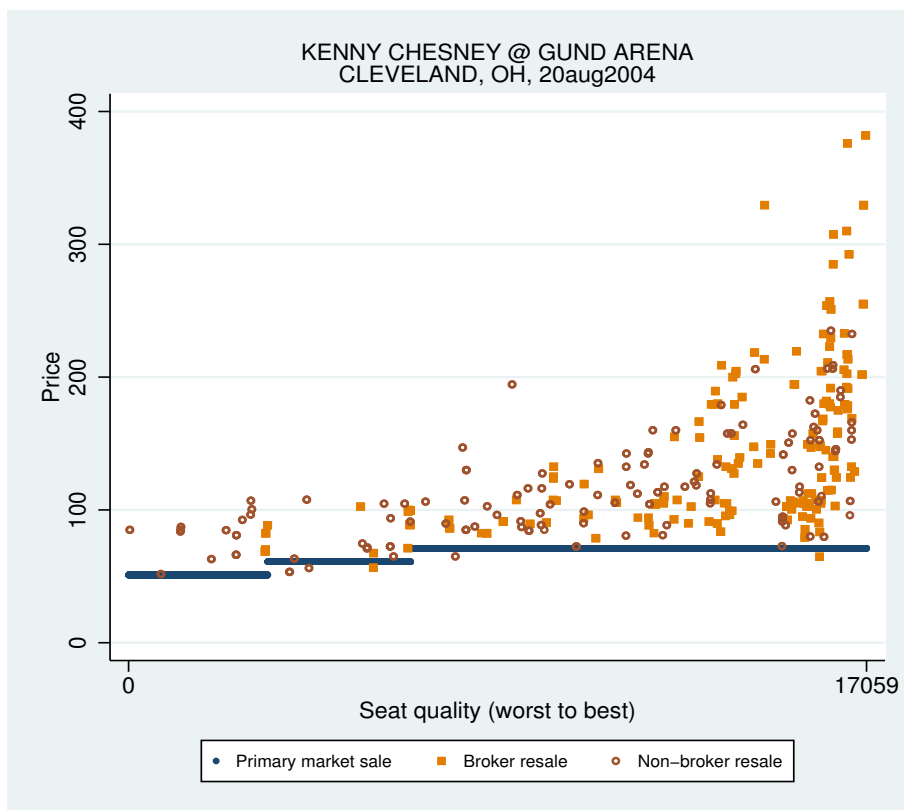
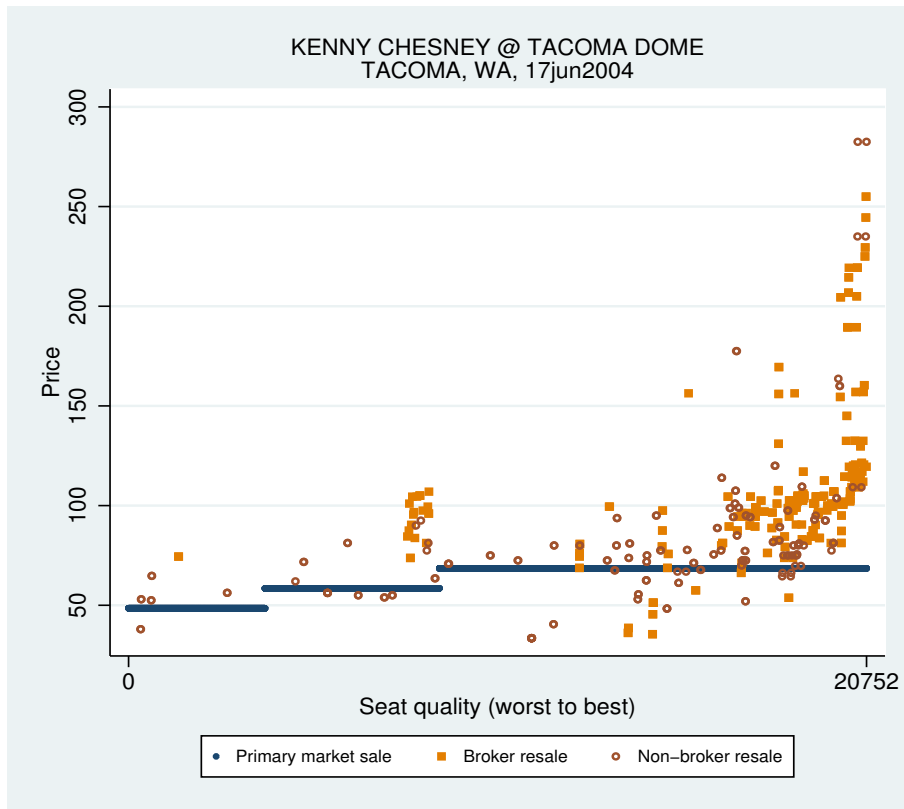


Figure 2: Two more sample events

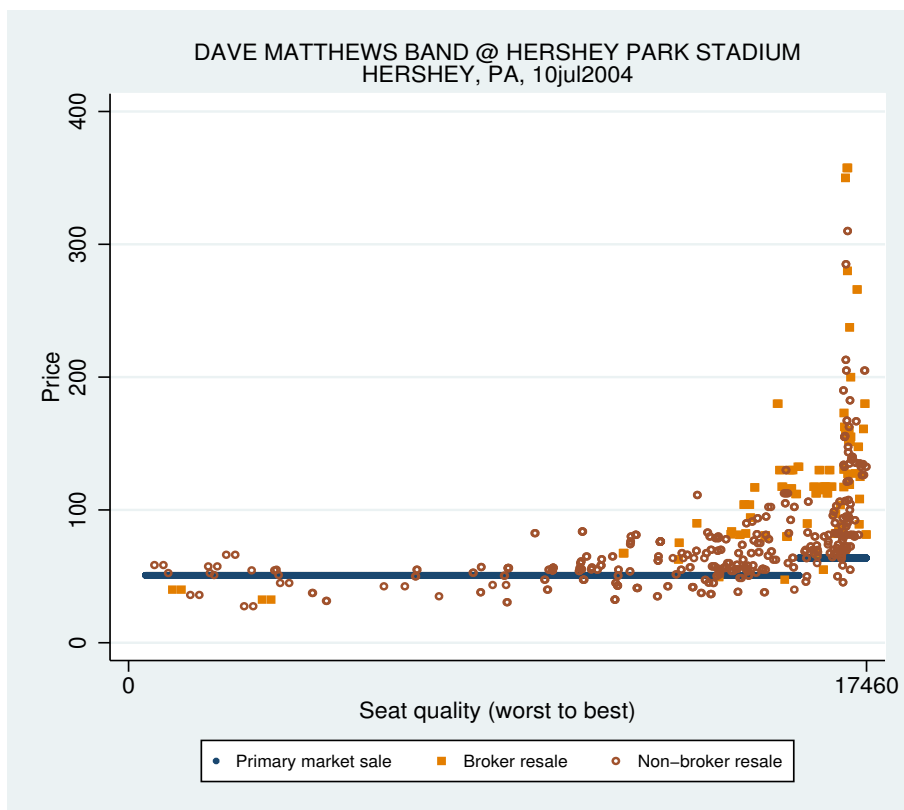



Figure 3: Probability of resale and relative seat quality




Figure 4: Resale prices and relative seat quality





Figure 5: Probability of resale by price level



In generating this figure, only events with three or more price levels were used. Relative seat qualities are calculated *within* price level for this figure, and the probability of resale is estimated using kernel-weighted local polynomial regression. So, for example, the probability of resale is on average higher for the best seats in price level 2 than for the worst seats in price level 1.

Figure 6: Timing of sales in primary and secondary markets



Time is normalized to make it comparable across events; it is measured as $(\text{days since onsale}) / (\text{total days between onsale and event})$. The histogram in the top panel represents the 1,739,346 tickets sold by Ticketmaster; the bottom panel represents the 68,828 tickets resold on eBay or StubHub.

Figure 7: The consumer's decision problem

