

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

Abstract

We develop an equilibrium model of ticket resale in which buyers' decisions in the primary market, including costly efforts to "arrive early" to buy underpriced tickets, are based on rational expectations of resale market outcomes. We estimate the parameters of the model using a novel dataset that combines transaction data from both the primary and secondary markets for a sample of major rock concerts. Our estimates indicate that while resale improves allocative efficiency, half of the welfare gain from reallocation is offset by increases in costly effort in the arrival game and transaction costs in the resale market.

October 2009

*Thanks to Lanier Benkard, Glenn Ellison, Michael Ostrovsky, and Marc Rysman for valuable suggestions. We are also 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.

1 Introduction

Ticket reselling is a controversial activity. In many jurisdictions (indeed in many countries) it is regulated or banned, and even where it is legal it is often stigmatized. Whether labelled as brokers, scalpers, or touts, ticket resellers are commonly loathed by concert artists, sports teams, and consumers. Roth (2007) even includes ticket scalping as an example of a “repugnant transaction.” This widespread hostility toward ticket resale seems at odds with the conventional view of economists, which is that voluntary trades made in resale markets result in more efficient allocations and increased social welfare.¹ Resale does not just affect final allocations, however. It also affects demand in the primary market, giving rise to initial allocations in the resale market that are endogenous to anticipated resale activity—resale causes profit-seeking individuals to purchase tickets who otherwise would not have, and causes some consumers to bypass the primary market and buy in the resale market. Moreover, by providing a profit-motive for ticket buyers, resale may stimulate individuals to engage in costly rent-seeking behavior as they compete to obtain the most valuable tickets in the primary market. For these reasons, not only does ticket resale give rise to redistributions of surplus in which some individuals are made worse off, but it may also reduce overall social welfare.

In this paper we present a new model of ticket resale, and we estimate its structural parameters using a unique dataset that merges transaction-level information from both the primary and secondary ticket markets for a sample of rock concerts. The model is designed to provide a comprehensive analysis of the welfare effects of reselling, including the efficiency gains from ticket reallocation, as well as changes in transactions costs incurred in the primary and secondary markets. There are three sequential stages of decision-making that we explicitly model. First, consumers make strategic choices of how much costly effort to exert in an arrival game that determines the sequence of buyers in the primary market. Second, individuals make choices about whether to buy in the primary market or to wait for the resale market. Third, purchased tickets are offered in the resale market to the endogenous pool of potential buyers that did not purchase in the primary market.

Importantly, buyers’ decisions in the primary market reflect rational expectations of resale outcomes. That is, buyers possess expectations of what will happen in the resale market, and their decisions in the arrival game and in the primary market lead to secondary market outcomes that are consistent (on average) with these expectations. Hence, for example, some types of buyers may choose to exert high effort in the arrival game driven by the expectation of reselling

¹See, Happel and Jennings (1995), Hassett (2008), McCloskey (1985), Mankiw (2007), and Williams (1994). In a 1999 *Forbes* magazine column, Dan Seligman asserted that anti-scalping laws are one of the ten dumbest ideas of the century (as referenced in Atkinson, 2004).

good tickets at high profit, and some types of buyers may choose to exert low effort in the arrival game because they expect to bypass the primary market and buy tickets from resellers. Incorporating forward-looking behavior is essential for analyzing the welfare consequences of resale because it explicitly allows the magnitude of transaction costs—and the allocation of tickets going into the resale market—to be endogenous to the degree of resale activity.

The amount of welfare gain that can result from reallocation depends largely on the degree of heterogeneity in consumers' valuations. In the model we allow consumers to vary in their willingness to pay for seat quality, as well as in their costs of exerting effort in the arrival game. We also allow these two dimensions of heterogeneity to be correlated (either positively or negatively), giving rise to the possibility that consumers who value good seats the most may tend to be either early or late in the arrival sequence. Even though consumers obtain utility from attending the event, the model includes the option for consumers to resell their tickets if they receive an offer exceeding their own reservation utility from using the ticket themselves.² The model also includes brokers, whose sole purpose is buying and reselling tickets to make a profit (they obtain zero utility from attending the event). Brokers have a potentially important impact on welfare since they provide additional liquidity to resale markets, helping to achieve an efficient final allocation, but they also capture some of the surplus for themselves. Brokers may also drive consumers to incur higher costs in the arrival game, as they stimulate competition for good tickets.

The key driver of resale activity in our model is underpricing in the primary market.³ Underpricing leads tickets to be inefficiently allocated in the primary market, leading to gains from trade that can be realized in the secondary market. In the data, underpricing is reflected mainly in the coarse price levels chosen by sellers. The rock concerts we analyze typically sell around 15,000 to 30,000 tickets, but these tickets are typically sold at only two different price levels. Hence, thousands of seats with significant variation in seat quality are all sold at the same price in the primary market, while in the resale market those seats are naturally resold at greatly differing prices. The conventional practice of coarse pricing in the primary market is puzzling, and we do not attempt to offer an explanation in this paper. We simply take primary market pricing as given, and focus our examination on the resale activity that arises as a result of it.

Our model has at least two novel implications. The first is that resale markets may reduce total welfare. When tickets are underpriced in the primary market, there is a strong profit incentive for individuals to buy tickets with the intent of reselling. This stimulates costly efforts

²In fact the offer must exceed the ticket holder's reservation utility plus the transaction cost in the resale market.

³Season ticket holders for sports such as baseball may also engage in resale as a form of unbundling. We study rock concerts where we are able to ignore this particular motivation for resale because rock concerts are rarely bundled.

to “arrive early” in the primary market and secure the best (most underpriced) seats. One might expect the incurred costs to be bounded above by the gains from reselling, but effort choices are strategic. The set of tickets available to an individual in the primary market depends on where the individual comes in the arrival sequence, which in turn depends on the effort choices of other individuals, and the resulting game is a prisoner’s dilemma. Consequently, while the presence of a resale market allows tickets to be reallocated in a way that increases total surplus, those efficiency gains may be more than offset by additional costs incurred in the arrival game.

A second implication of our model is that resale may reduce demand in the primary market. This result is surprising, because resale provides a valuable option to buyers in the primary market, and hence ought to increase primary market demand. However, a prohibition of resale leads to an inefficient allocation of tickets, and this inefficiency may actually cause more tickets to be sold. In the absence of resale, high-value consumers may end up with low-quality seats—and if high-value consumers value all seats more highly than low-value consumers, they may be the only consumers who would be willing to buy the low-quality seats. If, on the contrary, tickets could be costlessly resold, then all buyers would rationally anticipate that high-value consumers will end up with the high-quality seats, leaving no one who is willing to purchase the low-quality seats. Essentially, resale makes willingness-to-pay a steeper function of seat quality than it would be in the absence of resale, and the effect on primary market sales can more than offset the general increase in willingness-to-pay that comes from the option value of resale.

In order to assess the empirical relevance of these effects, we estimate the structural parameters of our model using very detailed data on a sample of 56 concerts. To our knowledge, this is the first study of ticket resale to utilize transaction data from both the primary market and the resale market. 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. The data contain information on individual ticket sales in the primary market, the speed of sales in the primary market (which is informative about the arrival game we analyze), which tickets are resold and at what prices, and whether the tickets were resold by brokers or non-brokers. Even from a purely descriptive standpoint, our analysis of these data constitutes a significant contribution to the understanding of ticket resale markets.

The data reveal several interesting facts about resale markets for rock concert tickets. While brokers account for the majority of resale activity, 46% of the resale transactions in our data are sold by non-brokers (i.e. consumers). On average, ticket prices in the resale market are 41% above face value. However, it is relatively common to see prices below face value: brokers (non-brokers) appear to lose money on 21% (31%) of the tickets they sell. The overall rate of resale is relatively low during our sample period, with only 5% of purchased tickets being resold on eBay or StubHub. Of course, for certain events this number is much higher. The event in our

dataset with the most active resale market had 17% of its tickets resold on eBay or StubHub, and resale market revenue was equal to 37% of the primary market revenue. The likelihood of resale is strongly associated with seat quality, as we find the best tickets are roughly four times more likely to be resold than low- to mid-quality tickets.

We find that observed levels of resale activity generate small improvements in allocative efficiency relative to a world without resale. However, these improvements come at a cost. Over half of the increase in gross surplus is offset by the combination of higher transaction costs in the resale market and higher costs of effort in the arrival game. Our estimates also imply that strategic behavior in the arrival game leads to primary market allocations that are significantly more efficient than a random allocation, as consumers with high willingness to pay try hardest to obtain tickets in the primary market. This highlights the importance of allowing for endogenous arrival—a model with random arrival would dramatically overstate the potential efficiency gains from resale.

We estimate that consumers have large transaction costs, preventing many exchanges that would otherwise be welfare-improving. Brokers' transaction costs are estimated to be much lower, and the results of our counterfactual analyses indicate that their participation in the market leads to a net welfare gain. In general, large reductions in transaction costs (for brokers and consumers) would lead to potentially significant increases in social efficiency. For example, we estimate that net social surplus (sellers' revenues plus buyers' net surplus) would increase by nearly 11% if resale markets were frictionless. We find that on average frictionless resale leads to higher demand in the primary market. However, for some events resale has the opposite effect, with primary market revenues falling by as much as 5% relative to a world without resale.

Finally, even though we estimate that resale increases aggregate surplus, our estimates show that not everyone is made better off. Under frictionless resale, for example, there is a large increase in surplus captured by ticket resellers, but a large decrease in the surplus earned by concertgoers. In other words, while resale reallocates tickets in a way that increases aggregate surplus, ticket resellers capture more surplus than they create. The biggest losers from resale are the consumers who actually attend the event.

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 discuss the expected welfare effects of ticket resale, and review the prior literature. In Section 4 we describe how we compiled the data, and provide summary statistics and descriptive analyses. The model is outlined in Section 5, and the details of the estimation, including identification, are described in Section 6. Section 7 discusses the results of various counterfactual simulations designed to assess the welfare consequences of resale, and Section 8 concludes.

2 Market overview

Live music and sporting events generates \$20 billion in primary market ticket sales in the U.S. each year. Reselling generates roughly \$3 billion in revenue each year, and this number is expected to grow over the next several years (Mulpuru and Hult, 2008). An important distinction from other ticketed products, such as airline travel, is that event tickets are usually transferable, which is necessary for legitimate resale activity. In this study we focus on music concerts, which allows us to avoid the complexity that season tickets, a major component of ticket sales for sports, would introduce to the analysis.⁴ Concerts are organized and financed by promoters, but the artists themselves are principally responsible for setting prices.⁵ Typically the artist and/or artist's 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 the logistics of ticket selling. The dominant firm in this industry is Ticketmaster, which serves as the primary market vendor for over half of the major concerts in North America. Ticketmaster sells tickets primarily online or by phone. Tickets usually go on sale three months before the event, and sometimes sell out on the first day.

Choosing primary market ticket prices is a complex problem for producers. Venues often have over 20,000 seats, with significant quality variation, implying many potential price-quality menus based on different partitions of the venue. Rosen and Rosenfield (1997) provide a theoretical analysis of how to divide a venue and what prices to set. TicketMaster has experimented with auctioning tickets in the primary market, but this is not yet common. The pricing problem may also be complicated by the possibility that artists have a preference for selling out the event, perhaps because artists obtain utility from playing to a full house, or because doing so enhances the experience for consumers.⁶ These considerations may explain why artists seem to routinely underprice their concerts. It is also conceivable that concert tickets are complementary to recorded music sales and other merchandise, in which case the objective is not to simply maximize ticket revenues. Artists sometimes also cite a desire to be fair or assure access for all fans.⁷ All of these are interesting issues, but in this study we do not seek to model or explain primary market pricing.⁸ In our empirical model we treat primary market prices as exogenous.

There is no federal regulation of ticket resale, but some states in the U.S. have laws forbidding or restricting ticket resale. For example, Arkansas, Kentucky, and Michigan prohibit reselling

⁴There may be instances of event bundling in rock concerts, but they are rare and do not apply to any of the events in our dataset. See Chu, Leslie and Sorensen (2009) for an analysis of ticket bundling.

⁵See Connolly and Krueger (2005) for a detailed review of the music industry.

⁶See Becker (1991), Busch and Curry (2005), and DeSerpa and Faith (1996).

⁷See Kahneman, Knetsch and Thaler (1986).

⁸See Courty and Pagliero (2009) for an empirical study of the determinants of when concerts utilize uniform pricing versus second degree price discrimination.

above face value. According to Fried (2004), as of 2004 there were 12 states with restrictions on resale, and 38 states without any restrictions (aside from limits on selling outside the venue). However, as best we can tell, these anti-scalping laws are rarely enforced by government authorities, perhaps because they are regarded as victimless crimes.⁹ Ironically, while eBay makes ticket reselling easier in many ways, eBay also enforces anti-scalping laws, so that these laws are perhaps more pertinent today than in the past. Elfenbein (2005) finds that there are fewer (but not zero) eBay transactions for NFL tickets in states with stricter regulations. There are various ways to get around the laws, even on eBay, such as charging exorbitant shipping fees. In any case, in the U.S. there is a clear trend of deregulation: in 2007, for example, Connecticut, Pennsylvania, Minnesota, Missouri and New York all repealed anti-scalping laws. Outside of the U.S., regulatory attitudes are also mixed. In the U.K., the Office of Fair Trading studied ticket reselling in 2004 and recommended against regulation.¹⁰ Scalping is illegal in some states in Australia, and was illegal for the 2008 Olympics in Beijing.¹¹

The internet has transformed the ticket resale industry. It is widely acknowledged by industry insiders and in the trade press that eBay is the dominant marketplace for ticket resale, followed by StubHub.¹² Two pieces of evidence support this belief. First, in a survey of concertgoers at a major rock concert in 2005, Alan Krueger found that eBay and StubHub accounted for between a third and a half of all resold tickets (see Connolly and Krueger, 2005). Second, based on survey evidence from 2007, Mulpuru and Hult (2008) report that eBay and StubHub account for 55% of online ticket resales. Tickets are also resold on numerous other web sites (Razorgator, TicketsNow, TicketLiquidator, etc), as well as offline.

Note that eBay and StubHub are marketplaces—unlike brokers, they do not own the tickets that are for sale. These web sites create value by lowering transaction costs. Features at some sites include interactive seating charts, personalizations, and parking and weather information. StubHub especially has a reputation for providing a well-designed and user-friendly interface. The web sites also extract value by charging fees on each transaction: eBay charges a transaction fee of about \$7 for a ticket that is resold for \$100, and StubHub charges \$25 for a ticket that is resold for \$100. To address fraud, eBay emphasizes their reputation mechanism, and StubHub provides a guarantee.

⁹In a study of resale laws, Elfenbein (2005) also notes that he is yet to find an instance of online ticket resale being prosecuted by an authority.

¹⁰Office of Fair Trading (2005).

¹¹The International Olympic Committee requires host countries to ban reselling.

¹²In January 2007, StubHub was acquired by eBay. We study data from before the acquisition.

3 Prior research into the welfare effects of ticket resale

We are certainly not the first to point out that primary market outcomes depend on buyers' expectations about the resale market. Several papers address the question: does resale increase or decrease producers' profits? Williams (1994) finds that NFL primary market prices in 1992 were lower in states with anti-scalping laws, suggesting that resale is good for producer profits. Depken (2007) studies more recent evidence and finds the reverse. Swofford (1999) provides the first theoretical analysis of this question, proposing that brokers have different risk attitudes, cost functions, or revenue functions than producers, which allows them to capture some consumer surplus without affecting producers' profits.¹³ Producers may nonetheless lobby for anti-scalping laws, since reselling causes consumers to pay higher prices without creating higher profits for producers.

Subsequent theory by Courty (2003) suggests producer profit may be invariant to resale activity. Courty's model incorporates demand uncertainty at the individual consumer level (i.e. consumers don't know their own value for the event until just before attending). The producer has the option of selling tickets at different prices on different dates, and has the ability to commit in advance. Courty shows that the producer's profit from selling all tickets at the last minute to informed consumers (with no resale possibility) is equal to the profit from selling all tickets early and allowing brokers to buy and resell tickets. Karp and Perloff (2005) present a model with asymmetric information in which broker activity conveys information to consumers about the value of tickets, allowing producers to charge higher prices. Geng, Wu and Whinston (2004) also argue that resale may be beneficial to producer profits.

A few papers address the question: does resale benefit or harm consumers? Thiel (1993) argues that the presence of brokers may be good for some consumers and bad for others. Thiel assumes excess demand in the primary market (with exogenously low prices). The presence of brokers reduces the probability of a given consumer obtaining a ticket in the primary market, but provides a second chance for consumers to obtain a ticket in the secondary market. Hence, brokers may benefit consumers that would otherwise miss out on a ticket in the primary market, but harm consumers that otherwise would have purchased a ticket at a lower price in the primary market. The net effect on consumer surplus is unclear. Busch and Curry (2005) argue that resale is harmful to consumers, based on a model with social externalities among consumers. On the empirical side, Elfenbein (2005) estimates the effect of anti-scalping laws on resale prices for NFL tickets.¹⁴ He finds that stricter regulations result in fewer tickets being resold, higher

¹³Based on a similar characterization of brokers, Spindler (2003) argues that brokers may enhance producers' profits: if brokers are better at extracting consumer surplus than producers, this may allow the producer to charge higher prices than without brokers.

¹⁴See also Hassett (2008).

resale prices and markups. This suggests an increase in resale activity results in lower resale prices, which may be attributable to competition among resellers. The implications for consumer surplus are unclear.¹⁵

Our approach differs from the prior research in four fundamental ways. First, unlike several of the prior studies mentioned above, we do not allow the producer to sell tickets in multiple periods. In the model we present below, the producer sells tickets in the primary market in period one, and resale activity takes place in period two. This is a simplifying assumption, motivated by the fact that there is minimal overlap in the timing of primary market activity and resale activity (as we show in the next section), and also by the fact that no producer in our dataset implements any form of dynamic pricing—primary market prices are equal over time for every event.¹⁶ But our approach does preclude us from evaluating any counterfactual pricing scheme in which producers vary primary market prices over time.

Second, we allow consumers (in addition to brokers) to resell tickets. This is important because it fundamentally affects how consumers value tickets in the primary market, and also because it allows consumers to capture some of the rents from reselling that only accrue to brokers in the prior research. Brokers may still be at an advantage over consumers in their role as resellers, because we allow for the possibility that brokers have lower transaction costs. Another benefit of our approach is that we can implement counterfactuals that preclude brokers from participating in the market while still allowing consumers to resell, or we can prohibit all resale activity. Perhaps even more importantly, our approach reflects reality—46% of the resold tickets in our dataset appear to have been sold by non-brokers (as reported in the data summary, below).

A third distinguishing feature of our approach is that we assess the welfare of consumers who attend the events separately from the welfare of consumers and brokers who engage in reselling. This is important because, as we show in our counterfactuals, resale activity can be harmful to the welfare of consumers who attend events, and beneficial to consumers who engage in resale. Moreover, the latter effect can outweigh the former—resale is good for consumer surplus, even though those who attend the events are made worse off. To the best of our knowledge, no prior paper has explicitly considered this distinction.

Fourth, we incorporate a strategic game of costly effort to determine the arrival sequence of buyers in the primary market. None of the prior papers incorporate into the analysis the impact of resale on costly efforts in the primary market. This is an important extension for a

¹⁵Some of these issues are also addressed in a recent study of resale price dynamics by Sweeting (2008).

¹⁶Note, however, that our model includes uncertainty of the same kind emphasized in Courty (2003). Namely, consumers do not know whether they will be able to attend the event in period one (due to the possibility of a schedule conflict).

couple of reasons. It allows us to take into account the effect of resale on primary market costs. It allows us to incorporate non-random arrival in the primary market—i.e. correlation between consumers’ valuations and their arrival order—which itself is based on a formal model of decision making. As we explain below, this feature also provides us with an estimation strategy that relies on less restrictive assumptions.

4 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 details about tickets sold through the primary market vendor (Ticketmaster), as well as information about all tickets that were resold on eBay and StubHub.

Two unique features of the dataset are especially valuable for this research. First, by merging data from Ticketmaster, eBay, and StubHub, we observe primary and resale activity in parallel. We suspect that data from either market in isolation would still be interesting and informative. But by combining data on primary market sales and resales, we are able to study the *interaction* between the two markets, which we believe is crucial to understanding how these markets work. Second, as we explain below, we distinguish resale activity by brokers from resale activity by consumers. For the reasons explained in Section 3, this is important to evaluating the welfare effects of resale.

Figure 1 illustrates the data for two example concerts: one by Kenny Chesney and another by the Dave Matthews Band. In each graph, the vertical axis represents price, and the horizontal axis represents seat quality, ordered from worst to best (we explain the measure of seat quality in more detail below). Consider the first panel, which shows the data for a Kenny Chesney concert performed in Tacoma, Washington. 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 (i.e. consumers) 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 Figure 1 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.

4.1 Primary market data

The primary market data were provided by Ticketmaster. The sample includes 56 concerts performed by 12 different artists: Dave Matthews Band, Eric Clapton, Jimmy Buffett, Josh Groban, Kenny Chesney, Madonna, Phish, Prince, Rush, Sarah McLachlan, Shania Twain, and Sting. In total there were 1,034,353 tickets sold or comped in the dataset. 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).¹⁷ 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. The information on the timing of sales in the primary market is crucial for our analysis of the arrival game, detailed below.

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 ranked above seats toward the outsides of rows further forward, and seats at the front of upper-level sections are sometimes ranked above seats at the back of lower-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

¹⁷A “comped” or complementary seat is one that was given away. Comps are typically around 1% of ticket sales (and are always less than 3%) for the events in our sample. An open seat is an available seat that went unsold.

total number of tickets available. The best seat ($j = 0$) therefore has quality 1, and the worst seat has quality $1/J$.

4.2 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 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 her 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). This information was provided directly by StubHub, and includes some details that are not available at their web site.

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.¹⁸ 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

¹⁸Dropping these sales means that our data will slightly understate the total amount of resale on eBay and StubHub for these events.

observe 4 tickets resold in section 101, unknown row, we assign those 5 tickets to seats 40, 80, 120, and 160 (of the 200).¹⁹

For the empirical model we estimate below, it would be ideal to observe *all* resale activity for the sample concerts. We do not know exactly how much of total resale activity is accounted for by eBay and StubHub. As explained in Section 2, eBay is almost certainly the largest single outlet for ticket resales, with StubHub likely the second largest. Where necessary in our analysis below, we simply assume that the combined market share of eBay and StubHub is 50%, and later test the sensitivity of our findings to this assumption. Note that even if we knew eBay’s and StubHub’s exact market shares for rock concerts in the summer of 2004, we would have no way to verify if resales on these two sites are representative of resale activity more broadly. However, 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 more broadly.

4.3 Summary statistics

As mentioned above, the dataset covers 1.03 million tickets sold in the primary market for 56 concerts by 12 different artists. Table 1 provides more detailed summary statistics. The capacity of each concert varies from 3,171 to 35,062 (the median is 17,483). The events were mostly sold out: over 75% of the concerts in our sample sold 100% of their capacity.

The average ticket price in the primary market was \$90.54 (including shipping and fees).²⁰ 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 \$144.15. There is also some within-event price variation, but not much. Most events tend to offer tickets at only two different price levels. The maximum number of price levels for a single event in our data is 4, and 6 concerts sold all tickets at a single price. The bottom panel of Figure 1 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.

We observe over 51,000 of the tickets purchased in the primary market being resold at eBay or StubHub (i.e. 4.96% of the number of tickets). As shown in Table 1, the maximum number

¹⁹We 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.

²⁰To be clear, \$90.54 is the average across events of the average ticket price at each event. If all tickets are weighted equally, the average ticket price in the primary market is \$81.04.

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 3% and 6%. On average, total revenue to resellers is equivalent to 8% 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 and StubHub alone, so they represent a lower bound for the total amount of resale activity.

Table 2 provides additional summary statistics of resale activity.²¹ The average resale price is \$113.23. Resellers make significant profits: the average markup is 41% over the face value, and 25% of resold tickets obtain markups above 67%. On the downside for resellers, 26% 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 \$90.43, while the average price of all purchased tickets in the primary market is \$81.04.

One reason why resold tickets tend to have relatively high face values is that resold tickets tend to be for relatively better seats. The average seat quality of tickets purchased in the primary market is 0.50, but in Table 2 we report the average seat quality of resold tickets is 0.61 (median is 0.65). Figure 2 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.²² It is clear that while tickets of all qualities are sometimes resold, higher-quality tickets are significantly more likely to 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. Meanwhile, resale prices reflect small differences in seat quality—every seat may have a different price. Figure 3 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 the 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.

Another important difference in how seat quality is priced in the primary and secondary

²¹In Table 1 an observation is an event, while in Table 2 an observation is a resold ticket.

²²In 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.

markets relates to monotonicity. Primary market prices are weakly monotonically increasing in seat quality for a given event. In contrast, the examples in Figure 1 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 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 4 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 4 shows exactly this pattern.

Figure 4 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 primary market, tickets typically go on sale 3 months before the event date. In Table 1 we report that (averaged across events) 70% 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 5 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 26 days of the event, and 25% 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.²³

²³The dynamics of resale prices are similar to those reported by Sweeting (2009) for Major League Baseball tickets, with one exception. Prices in the secondary market decline gradually as the event date approaches, but unlike for baseball tickets, prices for the events in our data tend to increase slightly in the two-week period just

The empirical model we estimate below allows consumers to invest in early arrival—i.e., to compete to be first in line to buy tickets in the primary market. Consumers’ incentives to do so depend on the degree to which the tickets are underpriced. In fact there is substantial variation across events in how compressed the sales are in time. For about 10% of concerts, more than 75% of the seats are sold in the very first day. But the median concert sells only 25% of capacity in the first day, and less than 75% in the first week. This suggests that people make costly efforts to show up early when excess demand is expected to be high: if it were costless to show up early, we would not expect to see any concerts with sales so spread out over time. Indeed, the concerts in our sample with the largest resale markups also tend to sell a higher fraction of tickets in the first day.

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 54% of resold tickets (as reported in Table 2). One of the more important facts in our data is that 46% of resold tickets are sold by non-brokers (i.e. consumers). In the table we also report that 21% of the tickets resold by brokers were at prices below the purchase price in the primary market. By comparison, 31% of the tickets resold by consumers were sold below face value. One possible explanation for this difference is that unlike brokers (who never intend to attend the event), consumers 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.²⁴ In the structural analysis we obtain a specific estimate of the rate of schedule conflicts (taking into account other potential reasons for resale).

The total profit (i.e. aggregate markup) obtained from ticket resale in our data is slightly over \$1.17 million. This is equivalent to 1.4% of total revenue in the primary market for these events. As a measure of “money left on the table” this suggests a fairly modest amount of forgone profit by firms in the primary market. This may be misleading, however, because modified pricing policies that capture some of this value may also be more efficient at extracting consumer surplus. We address this issue in the counterfactual analyses in Section 7. And this number could be considered a lower bound, since we do not observe all resale activity (as mentioned above, we believe we observe about half of all resale activity).

Lastly, we wondered if resale prices depend on the number of tickets grouped together. In particular, do pairs of tickets tend to sell for a higher price (per ticket) than single tickets? This would affect modeling assumptions in the next section. In an unreported regression, we regress $\log(\text{ResalePrice})$ on event dummies interacted with seat-quality deciles (i.e. flexible event and

before the event.

²⁴Another possible explanation is that brokers are better at identifying events with significant excess demand.

seat quality controls), and dummy variables for each of 1,...,5 tickets resold together. We found that the number of tickets has no significant effect on the resale price.

5 Model

An important driver of resale is arbitrage: profit-seeking behavior that takes advantage of underpricing (of particular seats, at least) in the primary market. Underpricing implies excess demand, requiring a mechanism for rationing tickets to buyers. We adopt a standard approach, in which buyers make purchase decisions in a sequence, with choice sets that are updated to reflect purchases made by buyers who came previously in the sequence.²⁵ The ordering of buyers is far from innocuous in this analysis, however. If we assumed that buyers were ordered from highest willingness to pay to lowest, this would yield an efficient allocation of tickets in the primary market, eliminating the principal motive for resale. Assuming that buyers are randomly ordered is more plausible, but precludes the possibility that high-value consumers may tend to arrive early in the sequence (because the benefits of being early are higher for these buyers than for low-value buyers), or that high-value consumers may tend to be late in the arrival sequence (because high-value buyers tend to have a high opportunity cost of time).

For these reasons it is essential to let the data reveal the arrival sequence. However, even this is not enough. Simply estimating the correlation between willingness to pay and arrival order in the primary market would not reveal how that correlation would be different if resale were banned or if resale were frictionless. In other words, just as the initial allocation of tickets in the resale market is endogenous, the arrival order of buyers in the primary market is also endogenous. How much effort people exert to buy tickets in the primary market depends on the existence and characteristics of the subsequent resale market.

We therefore propose a model of resale with three sequential stages of decision making. In the first stage buyers (consumers and brokers) make strategic effort choices in an arrival game that determines the ordering of buyers in the primary market. In the second stage buyers make purchase decisions in the primary market. In the third stage the resale market takes place. The equilibrium of the model is one in which all buyers behave optimally given their expectations about payoffs in subsequent stages, and their expectations are on average correct given that all agents in the model are behaving optimally.

In the following sections we outline the structure of the model. To keep the exposition

²⁵Leslie (2004) implements a similar approach. See Mortimer and Conlon (2007) for an alternative approach to demand estimation with stock-outs.

as simple as possible, we defer some of the details (such as functional form assumptions and simplifications made to reduce computational burden) until Section 6.1.

5.1 Primitives

There are two kinds 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.

Consumers are heterogenous in their willingness to pay (ω) for seat quality, and in their cost (θ) of arriving early in the primary market. These two dimensions of heterogeneity are jointly distributed with marginal density function $f_c(\omega, \theta)$. Brokers' costs of arriving early in the primary market are distributed with marginal density $f_b(\theta)$.

If a consumer attends the event, she obtains gross utility $U(\nu; \omega)$, where ν is the seat quality. Buyers' efforts to arrive early in the primary market (in order to secure higher quality tickets) are costly: we use $C(t; \theta)$ to denote the cost of arriving at time t for a buyer of type θ . Letting p denote the ticket price in the primary market, a consumer who purchases a ticket in the primary market and attends the event gets net utility equal to

$$U(\nu; \omega) - p - C(t; \theta) .$$

For a buyer who purchases in the primary market and then resells at price r in the secondary market, the net payoff is

$$r - p - C(t; \theta) - \tau ,$$

where τ is the transaction cost of reselling the ticket. The consumer who buys the resold ticket earns net utility

$$U(\nu; \omega) - r - C(t; \theta) .$$

In this case $C(t; \theta)$ may be zero, since the consumer may have chosen not to make any effort to buy a ticket in the primary market.

In essence, the objective of our empirical exercise is to use data on prices (p and r), quantities, and the timing of primary market sales (t) to estimate the distributions of buyer heterogeneity (f_c and f_b) and the parameters of the utility and arrival cost functions (U and C). Having recovered these primitives, we can then simulate market outcomes under various changes to the market environment (e.g., reductions in the transaction cost, τ , or increases in the sophistication of primary market pricing).

5.2 Arrival game

An agent's type is defined as the triple (b, ω, θ) , with $b = 1$ for brokers. In the arrival game agents have private information about their own types and have common knowledge of the distributions of types and the number of players. Strategies are defined as mappings from types to arrival times: $t \in \mathfrak{R}^+$. The higher is t , the later the agent arrives in the primary market. The arrival cost function, $C(t; \theta)$, is decreasing in t , and the marginal cost of arriving earlier is increasing in θ (i.e., $\partial^2 C / \partial t \partial \theta < 0$).

The benefit of an early arrival time (low t) is that earlier arrivals get to purchase higher quality tickets. But the ordering of agents in the primary market depends on the arrival times chosen by *all* agents: only relative arrival times matter. Since types are private information, upon choosing t an agent is still uncertain about her place in the eventual sequence. In other words, letting z denote the relative position in the buyer sequence, from the perspective of an individual buyer the mapping from t to z is stochastic.²⁶ Moreover, the payoff associated with relative position z is also uncertain, because it depends on choices made by buyers earlier in the sequence, and it depends on uncertain outcomes in the resale market (which we describe in more detail below). We denote the expected payoff associated with arrival position z as $V_0(z; b, \omega)$. The dependence on b and ω reflects the fact that expected payoffs differ for brokers versus consumers, and that (for consumers) payoffs depend on the marginal value of seat quality. Naturally, V_0 is decreasing in z for all buyers: early positions are the most valuable.

Agents in the arrival game therefore solve the following maximization problem:

$$\max_t \sum_z V_0(z, b, \omega) g(z|t) - C(t; \theta),$$

where

$$g(z|t) = H(t)^{zM-1} (1 - H(t))^{M-zM} \binom{M-1}{zM-1}.$$

In this notation $g(z|t)$ is the probability of being in relative position z given arrival time t , and the function $H(t)$ is the equilibrium distribution of chosen arrival times across all agents. Because the number of buyers in our empirical application is very large, for purposes of estimation we treat the mapping from t to z as deterministic, with $z = H(t)$.²⁷

The equilibrium distribution of arrival times can be obtained from an iterative procedure. Given V_0 and an initial conjecture $H_0(t)$, we calculate optimal arrival times for a large sample of buyers drawn from the distributions f_c and f_b . The distribution of these arrival times becomes

²⁶By "relative position" we mean that if a buyer is tenth out of one hundred, then $z = 0.1$ (for example).

²⁷For large M , $g(z|t)$ converges to a point mass on $E(z|t)$, which is just $H(t)$.

our new estimate, $H_1(t)$, and we then recalculate the arrival times that buyers would have chosen if they believed other buyers' arrivals were distributed according to $H_1(t)$. We iterate in this way until $\sup_t |H_s(t) - H_{s-1}(t)| < 0.01$. In practice, this procedure converges very quickly.

Importantly, note that the inclusion of an arrival game makes the welfare impact of resale ambiguous. Resale increases total surplus by reallocating tickets to consumers with the highest valuations, but it may also increase buyers' costly efforts in the arrival game—and these costs may more than offset the gains from reallocation. To understand why, consider a toy example in which a single ticket is sought by three potential consumers with net valuations of 30, 20, and 10. By incurring a cost of 6, each buyer can “arrive early.” If a buyer is the only one to arrive early, she gets the ticket with probability one; if two or more buyers arrive early (or if no one incurs the cost), each has an equal chance of getting the ticket. In the absence of resale, only the two consumers with the highest valuations will arrive early, and expected surplus is $\frac{1}{2}(30 + 20) - 2(6) = 13$. If the ticket can be costlessly resold, then its value becomes 30 for all three consumers, and all three have an incentive to incur the arrival cost. Expected surplus is $30 - 3(6) = 12$: the additional costs incurred in the arrival game (6) more than offset the gains from reallocation (5). The key idea is that the possibility of resale increases low-valuation consumers' (or brokers') incentives to obtain the tickets, increasing costly effort in the arrival game.

5.3 Primary Market

In the primary market stage, buyers make purchase decisions in the order that was determined in the arrival game. 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. When making their purchase decisions, buyers are forward looking. Consumers know that they will either consume the ticket (i.e., attend the event) or resell the ticket. Brokers who purchase in the primary market will always try to resell the ticket. We assume that both brokers and consumers incur transaction costs if they choose to resell, denoted τ^b and τ^c , respectively. The buyers in the secondary market are the consumers who chose not to purchase (or were rationed) in the primary market.

Buyers' decisions in the primary market are driven by their expectations about the resale market. Our model incorporates four sources of uncertainty about resale market outcomes. The first is randomness in the arrival sequence, 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 is the possibility of unanticipated schedule conflicts.²⁸ We assume there is a probability ψ that a given consumer will have zero utility from attending the event, with the uncertainty being resolved between the primary and secondary market stages. 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.

Randomness in auction participation is the third source of uncertainty. As explained below, we clear the secondary market using a sequence of auctions, with 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.

The fourth source of uncertainty is an aggregate (event-specific) shock to demand. We assume that the distribution of willingness to pay (f_c) is subject to shocks that are unobservable at the primary market stage. Buyers know the distribution of these shocks, but only observe the realized value of the shock after the primary market stage is complete. Incorporating this fourth 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, uncertainty about the strength of demand 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 uncertainties described above. For notational convenience, let Ψ be a random variable representing all four sources of uncertainty, and define the resale price function $R(\nu|\Psi)$ to be the resale price for a ticket with quality ν given the realization of Ψ —i.e., the arrival sequence, the schedule conflicts, the subsets of bidders participating in each resale auction, and the demand shock.²⁹

The decision problem for a broker in the primary market is to purchase the ticket j that maximizes

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

where p_j is the primary market price of ticket j , and $E(R(\nu_j|\Psi))$ is the expected price of ticket j in the resale market, where the expectation is with respect to the four sources of uncertainty

²⁸This source of uncertainty is equivalent to the uncertainty emphasized by Courty (2003) in his model of ticket resale.

²⁹The details of how we implement these features of the model are explained in Section 6.1 below.

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 for a payoff of zero.³⁰

A consumer's decision problem is somewhat more complicated, as illustrated in Figure 6. If a consumer buys ticket j in the primary market, with probability ψ she will be forced to resell the ticket, obtaining some price $R(\nu_j|\Psi)$. 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 $U(\nu; \omega) - p_j$. The expected payoff from buying ticket j is therefore

$$E(u^c|\text{buy } j) = -p_j + \psi E(\max\{0, R(\nu_j|\Psi) - \tau^c\}) + (1 - \psi) E(\max\{0, R(\nu_j|\Psi) - \tau^c, U(\nu; \omega)\}).$$

Again, the expectations are with respect to Ψ (arrival sequences, schedule conflicts, bidder participation, and demand shocks).

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

$$E(u^c|\text{wait}) = (1 - \psi) E(\max\{0, U(\nu; \omega) - R(\tilde{\nu}|\Psi)\}).$$

In this case, the consumer is not only uncertain about what the prices will be in the resale market, she is also uncertain about which ticket (if any) she will be able to buy. 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.

5.4 Resale market

The result of the primary market stage is an allocation of tickets to buyers. Some brokers and consumers hold tickets, and some consumers remain without tickets (either because they elected to wait for the secondary market, or because the event sold out before their turn in the buyer sequence). This allocation is not necessarily efficient, since the consumers without tickets may have higher willingness to pay than some ticketholders. In the resale market stage, ticketholders have the opportunity to sell their tickets to higher-value consumers. An important feature of our model is that the pool of potential bidders in the resale market is endogenous, consisting

³⁰We omit costs incurred in the arrival game from the present discussion, because those costs are already sunk when the primary market decision is made.

of consumers that either decided to wait and buy a ticket in the resale market, and consumers that were rationed out of the primary market due to the purchases of others.³¹

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. Although this approach is feasible in our model, it has one major drawback: it predicts resale prices that are monotonic in seat quality, which is very far from what we observe in the data. While resale prices increase on average as a function of seat quality, there is considerable variance in observed prices conditional on seat quality.

To accommodate this feature of the data, we clear the resale market using a sequence of private-values, second-price auctions with limited bidder participation.³² We begin with the highest quality ticket and 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.³³ If the offer is below the reservation price, the ticket remains with the seller. In this case, if the seller is a consumer, she uses the ticket herself and gets the consumption utility defined above; and 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. In this mechanism every ticket purchased in the primary market is for sale in the resale market, regardless of whether it is owned by a broker or consumer.³⁴

5.5 Equilibrium

Given the structure of payoffs described above, a rational expectations equilibrium is one in which: (i) brokers and consumers make decisions optimally in the arrival and primary market stages given their expectations about payoffs in the final stage (the resale market); and (ii) those

³¹All potential bidders must also have avoided a schedule conflict.

³²Assuming an auction mechanism in the resale market also corresponds with the actual functioning of this market.

³³We 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. We also rule out reselling any ticket twice. See Haile (2001) for an analysis of auctions followed by resale.

³⁴Note that even in the limit as L grows large, our approach differs from the “market-clearing price vector” approaches proposed for clearing assignment markets (e.g., Shapley and Shubik (1972) and Crawford and Knower (1981)). Instead of assuming that all buyers and all sellers are in the market at the same time and are aware of all other traders, we assume that buyers arrive randomly and sequentially, and do not anticipate participating in later auctions if they lose the current auction. Hence, while buyers in our model are forward-looking in the primary market stage, within the resale stage they behave myopically.

expectations are on average correct given optimal decision-making in the arrival game and primary market.³⁵ The challenge is finding expectations that rationalize a set of arrival times and 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 distribution of Ψ (or its components). Although realizations of this random variables lead deterministically to a set of resale market outcomes, the form of the function $R(\nu|\Psi)$ is not known. Nor is it possible to determine the value of $\tilde{\nu}$ as a function of Ψ .

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. We do this separately for expectations at the arrival game stage and the primary market stage, since the information set is slightly different at each of these stages. In particular, buyers in the arrival game are uncertain about which seats they will be able to buy in the primary market, because they cannot anticipate the exact purchase decisions of buyers who come ahead of them in the sequence. At the primary market stage, however, buyers know exactly which seats are available, and the only remaining uncertainty is about resale market outcomes.

Consider first the primary market stage. 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 $V_1(b, \nu, \omega|\gamma_1)$ to represent buyers' expectations at the primary market stage, where b is an indicator for whether the buyer is a broker, and γ_1 are the parameters.

The algorithm for finding a fixed point is as follows:

1. Choose an initial set of parameters, γ_1^0 . 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 $V_1(b, \nu, \omega|\gamma_1^0)$.
2. Use the *realized* final utilities from the simulations in step 1 to re-estimate the function $V_1(b, \nu, \omega|\gamma_1)$. Essentially, we regress realized utilities on a function of b , ν , and ω to obtain

³⁵Forward looking consumer behavior with rational expectations of future market outcomes is also essential in recent papers by Gowrisankaran and Rysman (2007), and Hartmann and Nair (2007). See also Chevalier and Goolsbee (2009).

a new set of parameters, γ_1^1 .

3. Use the new set of parameters from step 2 to simulate primary and secondary market outcomes as in step 1. Iterate on steps 1 and 2 until V_1 converges—i.e., until $V_1(b, \nu, \omega | \gamma_1^n)$ is sufficiently close to $V_1(b, \nu, \omega | \gamma_1^{n-1})$.

In the actual estimation, we use a very simple parameterization of V_1 . Letting h be an indicator for whether the buyer holds a ticket going into the second period, we let

$$\begin{aligned} V_1(b, \nu, \omega | \gamma_1) = & b \cdot h \cdot (\gamma_{10} + \gamma_{11}\nu) + (1 - b) \cdot h \cdot (\gamma_{12} + \gamma_{13}\nu + \gamma_{14}\omega + \gamma_{15}\nu\omega) \\ & + (1 - b) \cdot (1 - h) \cdot (\gamma_{16} + \gamma_{17}\omega). \end{aligned} \quad (1)$$

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 set of expectations V_1 such that the primary market choices that follow from V_1 lead to secondary market outcomes consistent with V_1 . The convergence criterion we use is based on average differences in V_1 . At each iteration of the algorithm, we essentially estimate the regression described in equation (1) using $M \times S$ “observations.” We stop iterating when

$$\frac{1}{MS} \sum_{i=1}^{MS} \left(\frac{|V_{1i}(\gamma_1^n) - V_{1i}(\gamma_1^{n-1})|}{V_{1i}(\gamma_1^{n-1})} \right) \leq 0.005.$$

In other words, we stop when the fitted values of V_1 differ from those of the previous iteration by less than half of one percent on average.

At the arrival game stage, buyers’ expectations about final payoffs are not a function of ν , because there is uncertainty about the seat qualities that will remain at the buyer’s turn in the sequence. We therefore approximate expectations as $V_0(z, b, \omega | \gamma_0)$, and use an iterative procedure analogous to the one described above to find a fixed point for V_0 . Namely, we begin with a conjectured set of parameters γ_0 , solve the arrival game given the implied V_0 , determine primary and secondary market outcomes given the resulting arrival sequence (including finding a fixed point for primary market expectations V_1), and then regress the final payoffs on a simple function of z (relative arrival position), b (a broker dummy), and ω (the buyer’s willingness-to-pay parameter) to obtain a new estimate of γ_0 . We iterate until the fitted values of V_0 differ from those of the previous iteration by less than half of one percent on average.

The specific parameterization we use for V_0 is

$$V_0(z, b, \omega | \gamma_0) = b(\gamma_{00} + \gamma_{01}z + \gamma_{02}z^2) + (1 - b)(\gamma_{03} + \gamma_{04}\omega + \gamma_{05}z + \gamma_{06}z^2 + \gamma_{05}\omega z) .$$

6 Estimation and results

Because the equilibrium of the model described above cannot be derived analytically, we estimate the model by simulated GMM. For a given set of model parameters, we can draw a sample of consumers and brokers randomly from f_c and f_b and compute the equilibrium outcomes in each stage of the model. By averaging over a large number of these simulations, we obtain predictions about arrival times, primary market sales, and secondary market sales (including resale prices) as a function of the model's 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.

Not surprisingly, this estimation approach is computationally intensive. The fact that the iteration on V_1 is nested within the iterations on V_0 makes the computational burden especially heavy. To make estimation feasible, we make several simplifying assumptions. These assumptions, along with the specific functional forms used for the utility function (U) and the arrival cost function (C), are described in the next section.

6.1 Details

We assume that the distributions of buyers' types (f_c and f_b) are lognormal. Specifically, we assume that for consumers

$$(\log \omega, \log \theta) \sim N \left(\begin{bmatrix} \mu_\omega \\ \mu_\theta \end{bmatrix}, \begin{bmatrix} \sigma_\omega^2 & \sigma_{\omega\theta} \\ \sigma_{\omega\theta} & \sigma_\theta^2 \end{bmatrix} \right) ,$$

and for brokers $\omega = 0$ with

$$\log \theta \sim N(\delta_b \mu_\theta, \sigma_\theta^2) .$$

Thus we estimate the means and variances of willingness to pay and arrival cost, and for consumers we estimate the correlation between willingness to pay and arrival cost type. The distribution of arrival cost types for brokers is assumed to have the same variance as for consumers, but the mean is scaled by δ_b . This allows for the possibility that brokers have better technologies for coming early in the buyer sequence (as is often alleged in the news media), in which case we would expect our estimate of δ_b to be less than one.

Let $\nu_j \in (0, 1]$ denote the quality of ticket j , measured as described in Section 4.1.³⁶ We assume consumer i 's gross utility from attending event k in seat j is

$$U_{ijk} = \alpha_k(1 + (\omega_i + \Delta_k)\nu_j^\phi), \quad (2)$$

where ω_i is consumer i 's willingness to pay for seat quality, and Δ_k is a mean-zero shock to the demand for event k . As explained above, event-specific demand shocks allow the model to explain why some sold-out events have low resale prices, while other events that didn't sell out can have high resale prices. Since buyers don't know the realization of Δ_k when they make their primary market decisions, there is some risk in purchasing tickets with the sole intention of reselling them. For purposes of estimation, we assume that $\Delta_k \sim N(0, \sigma_\Delta^2)$.

The chosen functional form for utility implies an intuitive interpretation of ω : when $\Delta_k = 0$, the ratio of a consumer's willingness to pay for the best seat ($\nu_j = 1$) versus the worst seat ($\nu_j = 0$) is just $1 + \omega$. The curvature term, ϕ , captures the potential nonlinearity of premia for high quality seats (as evidenced in Figure 3). 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 α_k .

Since estimating the α_k 's adds 56 parameters to an already difficult nonlinear optimization problem, we take a simple (but reasonable, we think) shortcut. In a preliminary step, we estimate event fixed effects in a regression of resale prices on seat quality and seat quality squared. We then assume that in the utility specification above $\alpha_k = \bar{\alpha} \cdot \hat{\alpha}_k$, where $\hat{\alpha}_k$ are the fixed effects estimated from the auxiliary regression. The $\hat{\alpha}_k$'s should effectively capture differences in the relative strength of demand across events. We then estimate only $\bar{\alpha}$, which allows the model to scale the event fixed effects up or down in order to improve the fit of the model.

As explained above, we clear the resale market with a sequence of auctions. We assume that the number of bidders in the auction for seat j is $L_j = 1 + \tilde{L}_j$, with $\tilde{L}_j \sim \text{Poisson}(\mu_L)$. We explain below how the data identify μ_L .

The parametric form of the arrival cost function is

$$C(t; \theta) = c_0 \left(\frac{\theta}{t} - 1 \right)^2, \text{ for } t \in (0, \theta].$$

Thus, if a consumer chooses $t = \theta$, she incurs no costs in the arrival game. The θ 's can be interpreted as the "exogenous" arrival times: the times at which buyers would have arrived in the primary market in the absence of any strategic efforts to arrive early.

³⁶If j is the ticket's position in the "best available" order, and there are a total of J available, then $\nu_j \equiv 1 - (j/J)$.

The arrival game is solved many times for every calculation of the GMM objective function, so it is a significant contributor to the overall computational burden of the estimation. To speed the computation, we make two simplifications. First, instead of solving for the equilibrium distribution of arrival times, $H(t)$, nonparametrically (which in principle is possible), we instead parameterize $H(t)$ as a Weibull distribution, and iterate on its two parameters. We found that the (converged) Weibull distribution is typically a good approximation to the actual distribution of arrival times. In other words, even though the actual distribution of arrival times chosen in equilibrium is not literally a Weibull, the Weibull provides a reasonably good fit.³⁷

As a final step toward making the computational 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.³⁸

To summarize, there are 15 parameters to be estimated: the parameters of the buyer-type distributions ($\mu_\omega, \mu_\theta, \sigma_\omega, \sigma_\theta, \sigma_{\omega\theta}, \beta, \delta_b$), the nonlinearity parameter in the utility function (ϕ), transaction costs for brokers and non-brokers (τ_b, τ_c), the probability of a schedule conflict (ψ), the standard deviation of the event-specific demand shock (σ_Δ), the mean number of bidders in the resale auctions (μ_L), the scaling parameter in the arrival cost function (c_0), and the scaling parameter for the event fixed effects ($\bar{\alpha}$). In the next section we explain which moments we use in the GMM estimation, and describe the variation in the data that identifies these 15 parameters.

6.2 Identification

A wide range of moment conditions could potentially be incorporated in the estimation. For the results reported below, we used a set of moments chosen to reflect the key sources of identifying variation in the data: the fraction of available tickets sold in the primary market (1 moment), average fraction of tickets resold by consumers (1 moment), average fraction of tickets resold by brokers (1 moment), average resale price (1 moment), average quality of resold tickets, separately

³⁷We also tried using Lognormal, Gamma, and Exponential distributions, but found that the Weibull assumption consistently yielded the best fit. The second simplification in the solution of the arrival game is that we discretize the set of possible arrival times. Each buyer i chooses t_i from a discrete grid on $(0, \theta_i]$. For the results reported below, this grid has 60 evenly spaced points.

³⁸This 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.

for broker resales and non-broker resales (2 moments), 25th and 75th percentiles of resale price distribution (2 moments), and 25th and 75th percentiles of resale seat quality distribution (2 moments), the fraction of primary market sales occurring in each of five time “buckets” (5 moments),³⁹ the fraction of first-day sales that are in the top price level (1 moment), the fraction of first-day sales that are in the second price level (1 moment), the fraction of sales in days 2-7 that are in the top price level (1 moment), the fraction of sales in days 2-7 that are in the second price level (1 moment), and the standard deviation of the residuals from a regression of resale prices on seat quality and seat quality squared (1 moment). Hence, we use a total of 20 moments to estimate the 15 parameters.

Formally, letting \tilde{s}_l and s_l denote the simulated and empirical values for each moment, we construct moment conditions of the form $m_l(\Theta) = \tilde{s}_l(\Theta) - s_l$, and select Θ to minimize $m'Wm$, where Θ is the set of all parameters, m is the stacked vector of moment conditions, and W is a weighting matrix. We employ the usual approach by first obtaining a consistent estimate of Θ , calculating an estimate of the optimal (variance-minimizing) weighting matrix, and then re-estimating Θ by minimizing $m'\hat{W}m$.

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 2.5 times the capacity of the event. The second is the fraction of total resales that our data account for. As explained above, we use the available information and 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).

For the estimated parameters we provide intuition for how basic patterns in the data provide identification. The “curvature” parameter ϕ is identified by the shape of the relationship between resale prices and seat quality (e.g., as shown in Figure 3). The shape of the price-quality relationship also influences the estimates of μ_ω and σ_ω , the mean and standard deviation of the distribution of $\log \omega$. However, these parameters are 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 we observe in the data that resale prices for the best seats are typically 3 times more than for the worst seats, then the distribution needs to be such that the highest draws of ω are around 2.

The amount of heterogeneity in willingness to pay (as captured by σ_ω) is also related to the observed variance in resale prices. In combination with μ_L , which determines the average

³⁹The five buckets are (day 1, days 2-7, days 7-14, days 15-30, days 31+).

number of bidders who randomly participate in each resale auction, this heterogeneity drives our model’s prediction of how “noisy” the relationship between resale prices and seat quality will be. If σ_ω is small and μ_L is large, for example, then our model would predict a very tight relationship. This identification argument is the rationale for including the standard deviation of the residuals from a regression of resale prices on seat quality as a moment to be matched in the estimation.

The standard deviation of demand shocks, σ_Δ , 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_Δ .

The fraction of buyers who are brokers (β) is mainly driven by the relative frequency of sales by brokers in the resale market. To be clear, however, the estimate will not simply equal the relative frequency of broker sales 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.

The probability of schedule conflicts, ψ , is driven by the relative rate at which consumers versus 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).

Identification of the transaction costs is driven by ticketholders’ relative propensity to resell at high versus 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. More specifically, the transaction costs estimates should depend on the slope of the relationship between the probability of resale and the expected markup, and particularly 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).

The arrival cost parameters (μ_θ , σ_θ , and c_0) are identified by the timing of purchases in the primary market. Intuitively, the data reveal the marginal benefit of accelerating arrival in the primary market. To the extent that resale prices capture the tickets’ market value, they also tell us how much more valuable it was to be 1st in the buyer sequence as opposed to 101st, say. For

events that were dramatically underpriced, this difference tends to be large, so we expect buyers to hurry and tickets to sell out very quickly. By contrast, for an event that is not underpriced, the incentives to arrive early are much weaker: only consumers with high willingness to pay for quality (high ω) have much incentive to hurry to the front of the line. Essentially, for any given event the data tell us (a) how valuable it was to come early and (b) how quickly the tickets sold (i.e., how hard buyers tried to come early), and observing this relationship across several events allows us to back out what the costs of early arrival must look like. Naturally, the moments related to the timing of primary market sales are intended to leverage this source of variation.

The estimate of δ_b , which represents the degree to which brokers' arrival cost distribution differs from the distribution for consumers, is driven by the difference in the average quality of tickets resold by brokers vs. consumers. In the data, brokers resales tend to be for higher-quality tickets, suggesting that they may be better than the average non-broker at "arriving early" in the primary market.

Finally, the correlation between arrival cost types (θ) and willingness to pay (ω) is identified by the relative demands for high-quality versus low-quality tickets early in the arrival sequence. Consumers who arrive early in the primary market can typically choose between high-quality, high-price seats and lower-quality, lower-price seats. If early arrivers tend to buy lower-quality seats (e.g., seats in the second price level or lower), this would suggest that θ and ω are positively correlated. This idea motivates the inclusion of moments measuring high- vs. low-quality sales in the early portion of the on-sale period. In addition to this source of identification, however, it is worth pointing out that the overall level of resale activity also influences the estimate of $\sigma_{\omega\theta}$. In the model, if ω and θ are negatively correlated, then the early arrivers also tend to be the buyers with the highest willingness to pay. In that case the primary market allocation is relatively efficient, leading to smaller gains from reallocation and fewer resales. If instead ω and θ are positively correlated, then the primary market allocation is very inefficient and (all else equal) we would expect to see a high volume of resale activity as tickets are traded from low- ω consumers to high- ω consumers.

6.3 Estimation Results

The estimates are reported in Table 3. Consumers' transaction costs are estimated to be about \$61, which may seem high. However, many consumers have never used eBay before and perceive there to be significant setup costs involved in doing so for the first time. Another interpretation is that the transaction cost captures an endowment effect (see Kahneman, Knetsch, and Thaler,

1990)—consumers’ valuations of tickets increase after purchasing them.⁴⁰ Brokers are estimated to have lower transaction costs than consumers, which is not surprising. Further, the mean of the arrival cost distribution for brokers is estimated to be 0.17 times the mean for non-brokers, suggesting that brokers have significantly better technologies for arriving early in the primary market. The estimated fraction of brokers (β) is 0.004, implying that there is one broker for every 250 consumers.

We estimate that consumers’ willingness-to-pay and arrival cost parameters (ω and θ) are negatively correlated. Thus, consumers who value the tickets highly will also tend to come earlier in the primary market buyer sequence. An implication is that the primary market allocation will be somewhat more efficient than a random allocation. We examine this issue more closely in the next section. The estimated distribution of ω is such that the average consumer is willing to pay 2.3 times more for the best seat than she is for the worst seat. A consumer at the 90th percentile of the distribution would be willing to pay about 3.8 times more.

The parameters of the arrival cost distribution imply that if no effort were exerted in the arrival game, only about 25% of the buyers would arrive in the first week of the onsale period. To arrive on the first day, a consumer at the 50th percentile of the distribution of θ would incur costs of roughly \$24, while a consumer at the 90th percentile would incur the same cost to arrive on day 7.

The estimate of μ_L implies that resale auctions have on average only 3.4 buyers participating. This number seems small relative to the number of buyers we observe submitting bids on eBay auctions, but not all of the bids we observe on eBay should be considered serious bids (many auctions begin with a number of lowball bids). In any case, this estimate is driven by the relatively high variance of resale prices (conditional on seat quality) that we observe in the data.

In general, the model fits the resale-related moments fairly well: the predicted rate of resale is very slightly lower than the rate observed in the data, and the predicted quantiles of resale prices are slightly higher than those observed in the data. The model underpredicts the average fraction of tickets sold in the primary market by nearly ten percentage points, and also underpredicts the fraction of tickets that are sold in the first day. We suspect the model would achieve a significantly better fit of the primary market moments if it were computationally tractable to estimate event fixed effects (i.e., the α_k ’s) directly.

⁴⁰Krueger (2001) argues that endowment effects are in fact important in ticket markets.

7 Counterfactual analyses

We now turn to our primary objective of quantifying the resale market’s impact on aggregate social welfare and the distribution of surplus among primary market sellers, brokers, and consumers. We do this by means of counterfactual analyses: i.e., given our estimates of the structural parameters, as reported in Table 3, we simulate market outcomes under various hypothetical changes to the market environment. It is important to note that we do not re-optimize primary market prices under the various counterfactuals, although of course we do incorporate endogenous primary and secondary market decisions of consumers and brokers.⁴¹

Table 4 compares outcomes under varying levels of resale frictions. To construct the table, we calculate averages across 100 simulated outcomes for each event (i.e., outcomes for 100 separate draws of Ψ), and then report averages across the 56 events. For the “base case” we simulate the model at the estimated parameter values. Outcomes in the no-resale case are simulated by setting the transaction costs (τ_b and τ_c) to arbitrarily high levels. To simulate outcomes with frictionless resale, we set transaction costs to zero *and* increase the number of participants in the resale auctions to essentially include all potential buyers (by increasing μ_L to a very high number).

Welfare consequences of resale

The first three columns of Table 4 are based on the estimated model, in which the buyer sequence is endogenous (with buyers playing a strategic arrival game). In the top row we report the *gross* surplus of the consumers who attend the event. The principal consequence of resale markets is to reallocate products to consumers with higher willingness to pay, and changes in the gross surplus of attendees capture the efficiency gains from this reallocation. To facilitate comparisons across regimes, we normalize all numbers in the table so that gross surplus equals 100 in the no resale case. There is no ex-ante ambiguity about the effect of resale on gross surplus of attendees—resale helps tickets end up in the hands of high value consumers. In terms of magnitude, we find that the actual level of resale in the data results in 1% higher gross surplus than if there was no resale. Under frictionless resale, gross surplus is 11.7% higher than the no resale case. Although not shown in the table, under frictionless resale 46% of tickets sold in the primary market are resold (on average).

However, there are costs to achieving this improvement in allocative efficiency. In the base case we find that the combination of transaction costs in the resale market and increased in

⁴¹While it is possible that sellers would adjust ticket prices in response to the hypothetical changes we consider, we do not know *how* prices would adjust, since it is not even clear that the sellers’ objective is to maximize profits.

costs of effort in the arrival game amounts to 60% of the gross surplus gain. Hence, while gross surplus increases by 1%, net surplus increases by only 0.5% (under the base case relative to no resale). Under the base case, arrival costs increase by a only a small amount. By comparison, in the frictionless resale regime we find that arrival costs significantly increase, because the possibility of costless resale increases buyers' incentives to compete for the best tickets. Under purely frictionless resale, *everyone* values the best seat at the willingness-to-pay of the highest- ω consumer.

We noted above that in principle the increased costs associated with resale could more than offset the improvement in allocative efficiency. While this is not the case on average for our estimated model, our simulations indicate that for some events net surplus would increase if resale were eliminated. That is, while gross surplus is always higher in the base case versus the no-resale case, for some events this increase is outweighed by the combination of transaction costs and increased arrival costs in the base case. At observed levels of resale activity, therefore, the impact of resale on net social surplus may be positive on average, but it is a close call. By contrast, we find no ambiguity in the frictionless resale case: relative to no resale, we estimate that frictionless resale would increase net surplus for all events in our sample.

The bottom panel of the table shows how the surplus is divided among the various market participants. Two findings are especially noteworthy. First, primary market revenues are higher (on average) in the presence of frictionless resale. As explained above, the effect of resale on primary market sales is theoretically ambiguous. In fact, the simulations indicate that for 9 of the 56 events the inefficiency of the no-resale allocation actually leads to more primary market sales, since high-value buyers end up consuming low-quality seats that otherwise wouldn't be purchased.

A second noteworthy finding is that resale significantly reduces the net surplus of event attendees. Under the base case, attendees' net surplus is 1% lower than the no resale case. Under frictionless resale attendees are 17% worse off. Interestingly, consumers that engage in reselling (non-broker in the table) are the biggest winners (under frictionless resale). This raises an interesting distinction: frictionless resale leads to a 15% increase in total consumer surplus, but a 17% decrease in surplus of consumers that actually attend the event. This is a symptom of the more general conclusion that reselling results in a transfer of surplus from attendees to resellers (which in our analysis happens to include consumers). Put simply, in a world with frictionless resale, consumers get the "right" tickets, but they pay a much higher price for them.

Importance of endogenous arrival

A central point of this study is that it is essential to model the impact of resale on primary

market behavior, in order to fully assess the consequences of resale activity. In particular, the efforts of buyers to obtain tickets in the primary market—which determine the allocation of tickets in the primary market—depend on whether resale is possible. In the last two columns of Table 4 we show how different our conclusions would be if we instead assumed that buyers arrive in a purely random sequence (regardless of whether resale is possible). That is, instead of allowing buyers to choose their arrivals strategically, we simply assign the sequence randomly in a way that is independent of willingness to pay.

Without resale, a purely random buyer sequence leads to an allocation that is 31.5% less efficient than the allocation that results from endogenous arrival. This is because endogenous primary market allocations are significantly more efficient than random allocations: high value buyers tend to invest in early arrival. Importantly, this also means that if we had estimated the model without endogenizing the arrival sequence, we would have dramatically overstated the potential gains from reallocation through resale. Under random arrival, frictionless resale leads to a massive increase in net surplus from 68.5 to 108.5—an increase of 58% (compared to 11% under the model with endogenous arrival).⁴² Attendees also benefit from resale under random arrival, contrary to the endogenous arrival model.

Strategic interaction in arrival game

As noted above, effort choices in the arrival game are strategic decisions because the position of any individual in the arrival sequence depends on the effort levels of other buyers (in addition to their own effort). Figure 7 graphically shows the strategic effect that resale has on arrival costs, based on the estimated model. The figure is based on averages across all events. On the horizontal axis are the deciles of the marginal (estimated) distribution of consumers’ willingness to pay for seat quality (ω). The solid line represents the percent change in average arrival cost for each decile of consumers, due to a change from no resale to frictionless resale. The dashed lined indicates the percent change in average position in the arrival sequence, for each decile.

It is evident from Figure 7 that low- ω consumers dramatically increase their arrival costs when resale is allowed, for the reasons explained above. Of greater interest is the fact that average arrival costs also increase for high- ω types. On the one hand, resale may cause some high- ω consumers to reduce their effort in the arrival game, preferring instead to let others incur those costs, and knowing that they have the option of waiting to buy a ticket in the resale market. On the other hand, if they do wish to purchase a ticket in the primary market—preferably an underpriced, high-quality ticket—then the high- ω types will need to increase their arrival efforts

⁴²Note that if resale were literally frictionless, then the gross surplus of attendees would be equal in columns 3 and 5 of the table. However, buyers’ inability to trade up in the resale market means that the final allocation still depends to some extent on the order of arrivals in the primary market.

as a strategic response to the higher efforts of the low- ω types. The figure indicates that the latter effect tends to outweigh the former.

The dashed line in Figure 7 shows that the (sometimes dramatic) increase in arrival effort results in barely any change in the arrival sequence. Recall that we estimate a negative correlation between ω and θ : high value consumers tend to have a low cost of effort. Combined with high- ω types' stronger incentives to obtain the best tickets, this leads to a no-resale equilibrium in which high- ω consumers tend to be early in the arrival sequence. As shown in the figure, frictionless resale causes these same consumers to increase their efforts in order to preserve their early position, further illustrating the importance of strategic interaction in the arrival game.

Role of brokers

Since many legal restrictions on ticket resale seem to be motivated by hostility toward brokers/scalpers/touts, in Table 5 we explore counterfactuals with variation in the presence of brokers. The first column reports results from simulating the model with β , the share of buyers who are brokers, set to zero. In the second column β is set to its estimated value of 0.004, and in the last two columns we experiment with significant increases in the presence of brokers, by increasing β to 0.02 and then 0.10. We normalize all values in the table based on the gross surplus under the base case (set to 100).

The table shows that increasing the presence of brokers leads to higher levels of gross surplus. This is because brokers provide liquidity to the resale market by virtue of having lower transaction costs than consumers. Increasing broker presence leads to a slight reduction in arrival costs. This is interesting because the frictionless resale counterfactual (reported in Table 4) revealed an increase in arrival costs. The reason for the difference is intuitive—frictionless resale increases the value of arriving early for low- ω consumer types, leading them to exert more effort; but increasing the presence of brokers has no impact on the value that any consumer places on any ticket, leading to no direct impact on effort choices. There remains an indirect effect on effort levels via strategic interaction, but this effect is apparently swamped by the simple fact that brokers themselves are estimated to have lower costs of arrival. Hence, it is not surprising that increasing the presence of brokers has a positive effect on net surplus.

As shown in Table 5, we find the effect of brokers on attendees' net surplus is less clear. A small increase in the fraction of brokers in the market is good for consumers, because the benefits from reallocation outweigh the higher prices paid to brokers. However, in our counterfactual with $\beta = 0.10$ we find that attendees are worse off than at any lower level of brokers (including no brokers at all). Notably, this is not because brokers capture more value than they create—indeed, comparing the counterfactuals of $\beta = 0$ with $\beta = 0.10$ we find that brokers capture 23%

of the value they create. The big winner from increasing the presence of brokers is the primary market seller.

Re-pricing best seats

We noted above that much of the observed resale activity in our data appears to be driven by unpriced seat quality. In particular, consumers evidently are willing to pay significant price premiums for the very best seats, but these seats are typically sold together with many inferior seats at the same coarsely-defined price level. To understand what would happen to resale activity if the best seats were re-priced, we simulated a separate counterfactual in which we took the top 10% of each event’s seats and assigned them a new price equal to the average observed resale price for those seats, and then simulated market outcomes using the parameter estimates from Table 3. Hence, under this counterfactual we add one additional price level to every event.⁴³ We find that primary market revenue increases by 13% (an average of approximately \$150,000 per event). This is a lower bound for how much money producers are leaving on the table by not scaling the house more finely. Setting higher prices for premium tickets also weakens buyers’ incentives to invest in early arrival: relative to the baseline model, arrival costs decline by 14%. These results reinforce an obvious but important point—the inefficiencies that give rise to resale markets could be mostly eliminated through improvements in primary market pricing.

Forward-looking versus myopic consumers

Finally, we also examined what happens if we ignore the interplay between primary and secondary markets—i.e., if we assume buyers are not forward-looking when they make their primary market decisions. We eliminate brokers, assign the buyer sequence randomly (i.e., no arrival game), and determine the primary market allocation by having consumers’ first-period decisions depend only on their consumption utilities from attending the event. We then clear the resale market as in the baseline model. Ignoring the potential profits from resale obviously reduces the perceived value of the tickets; our model predicts that primary market sales would decline by 1.2% if first-period decisions were made myopically.⁴⁴ The magnitude of the decline gives a rough indication of how many primary market purchases are speculative in the baseline model—i.e., how many tickets are purchased only because the buyer expects to resell for a profit.

⁴³In the other counterfactuals we consider changes to the fundamental parameters. In this case we kept the same parameters but considered a change to the data.

⁴⁴Recall that the level of resale activity in the base case is somewhat low. If transaction costs were low, this exercise would undoubtedly reveal a much larger difference between forward-looking consumers and myopic consumers.

8 Conclusion

A common complaint from consumers is that resale markets make it more difficult to obtain tickets in the primary market. However, before the internet boosted ticket reselling (by lowering resale transaction costs), consumers complained about the difficulty of getting tickets to popular events at all. Our modeling approach captures both of these effects. Resale stimulates competition for tickets in the primary market, making it costlier (in an effort sense) to buy in the primary market. But resale also makes it easier for consumers to buy tickets to any event in the resale market, as long as they are willing to pay market-driven prices. In other words, resale exacerbates the problems associated with excess demand in the primary market (i.e. costly rent-seeking behavior), but makes the final allocation of goods to consumers more efficient. This paper has sought to clarify these effects and empirically quantify their magnitudes.

Our approach has focused on the interdependence of primary and secondary markets, and is the first (to our knowledge) to analyze data from both markets in parallel. Our findings show that while the basic economics of resale markets are simple (buy low, sell high), the welfare consequences of resale—in particular, the *distribution* of gains and losses—are more subtle. In the market for rock concerts, we find that observed levels of resale activity do not generate dramatic welfare gains relative to a world without resale. However, substantial increases in social surplus could be realized by eliminating or reducing frictions in the resale market (e.g., transaction costs). To the extent that online marketplaces like eBay, StubHub, craigslist, and others facilitate secondary market exchanges by lowering transaction costs, we can infer that their services increase the total surplus generated by the market for event tickets.

Resale leads to a more efficient allocation of tickets, but does so at a cost. By enabling profitable resale transactions, it motivates individuals to engage in costly rent-seeking behavior in the primary market. Our analysis emphasizes how strategic interactions amplify these costs. We find that these costs are substantial. Comparing the observed level of resale to a counterfactual world with no resale, half of the gain in gross surplus from reallocation is offset by increased arrival and transaction costs.

Not everyone benefits from resale. In particular, consumers who attend the event may be worse off when resale markets become more fluid. Seats are allocated more efficiently—high quality seats end up being occupied by consumers with the highest willingness to pay—but the additional surplus generated by the improved allocation is mostly captured by resellers. As a group, concert attendees would have preferred less efficiently allocated tickets obtained at lower prices. We find that frictionless resale markets would lower the surplus of concert attendees by 17% on average. Thus, if the aim of public policy is to maximize *total* surplus (as arguably it

should be), then our findings provide some support for the repeal of anti-scalping laws. From a consumer protection standpoint, however, the conclusion may be different: if the narrow goal is to maximize the surplus of those who ultimately attend the event, then restrictions on resale may be warranted.

Our results also imply that resale markets on average lead to higher revenues for primary market sellers. More generally, our analysis highlights the degree to which resale activity is driven by pricing practices in the primary market. As primary market sellers implement increasingly sophisticated pricing schemes, resale activity should be expected to decline. Ticketmaster, for example, has actively encouraged artists to sell concert tickets using an auction mechanism. Indeed, one interpretation of why scalping arises is that brokers are more efficient at implementing such schemes (see Swofford, 1999), and as primary market sellers develop these capabilities themselves the value of brokers will diminish.

References

- Atkinson, J.J. (2004): “The Economics of Ticket Scalping,” *Mimeo*.
- Becker, G.S. (1991): “A Note on Restaurant Pricing and Other Examples of Social Influences on Price,” *Journal of Political Economy*, 99, 1109–16.
- Busch, L.A. and P. Curry (2005): “Rock Concert Pricing and Anti-Scalping Laws: Selling to an Input,” *Mimeo*.
- Chevalier, J.A. and A. Goolsbee (2009): “Are Durable Goods Consumers Forward Looking? Evidence from College Textbooks,” forthcoming in *Quarterly Journal of Economics*.
- Chu, S., P. Leslie and A. Sorensen (2009): “Bundle-Size Pricing as an Approximation to Mixed Bundling,” forthcoming in *American Economic Review*.
- Connolly, M. and A.B. Krueger (2005): “Rockonomics: The Economics of Popular Music,” *Mimeo*, Princeton University.
- Courty, P. (2003): “Ticket Pricing Under Demand Uncertainty,” *Journal of Law and Economics*, 46, 627–52.
- Courty, P. and M. Pagliero (2009): “Price Discrimination in the Concert Industry,” *Mimeo*.
- Crawford, V.P. and E.M. Knower (1981): “Job Matching with Heterogeneous Firms and Workers,” *Econometrica*, 49, 437–450.
- Depken, C.A. (2007): “Another Look at Anti-Scalping Laws: Theory and Evidence,” *Public Choice*, 130(1–2), 55–77.
- 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.
- Fried, J. (2004): “Admit Two: StubHub’s Founders Want to Take the Worry Out of Getting Close Seats,” *Stanford Magazine*, Nov/Dec 2004.
- Geng, X., R. Wu, and A.B. Whinston (2007): “Profiting from Partial Allowance of Ticket Resale,” *Journal of Marketing*, 71, 184–195.
- Gowrisankaran, G. and M. Rysman (2007): “Dynamics of Consumer Demand for New Durable Goods,” *Mimeo*.
- Haile, P. (2001): “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales,” *American Economic Review*, 91(3), 399–427.

- Happel, S.K. and M.M. Jennings (1995): “The Folly of Antiscalping Laws,” *The Cato Journal*, Spring.
- Hartmann, W.R. and H.S. Nair (2007): “Retail Competition and the Dynamics of Consumer Demand for Tied Goods,” *Mimeo*.
- Hassett, K.A. (2008): “Estimating the Consumer Benefits of Online Trading,” *Mimeo*.
- Karp, L. and J. Perloff (2005): “When Promoters Like Scalpers,” *Journal of Economics and Management Strategy*, 14(2), 477–508.
- Kahneman, D., J. Knetsch and R. Thaler (1986): “Fairness as a Constraint on Profit Seeking: Entitlements in the Market,” *American Economic Review*, 76(4), 728–41.
- Kahneman, D., J. Knetsch and R. Thaler (1990): “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy*, 98(6), 1325–48.
- Krueger, A.B. (2001): “Supply and Demand: An Economist Goes to the Super Bowl,” *Milken Institute Review: A Journal of Economic Policy*, 3(2), 22–9.
- Leslie, P. (2004): “Price Discrimination in Broadway Theatre,” *RAND Journal of Economics*, 35(3), 520–41.
- McCloskey, D. (1985): *The Applied Theory of Price*, 2nd Ed., New York: McMillan.
- McFadden, D. (1989): “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration,” *Econometrica*, 57, 995–1026.
- Mankiw, G.N. (2007): *Principles of Microeconomics*, 4th Ed., Thoms South-Western.
- Mortimer, J. and C. Conlon (2007): “Demand Estimation Under Incomplete Product Availability,” *Mimeo*.
- Mulpuru, S. and P. Hult (2008): *The Future of Online Secondary Ticketing*, Forrester Research Inc.
- Office of Fair Trading (2005): *Ticket Agents in the UK*.
- Rosen, S. and A. Rosenfield (1997): “Ticket Pricing,” *Journal of Law and Economics*, 40(2), 351–76.
- Shapley, L. S. and M. Shubik (1972): “The Assignment Game I: The Core,” *International Journal of Game Theory*, vol. 1, 111–130.
- Spindler, Z.A. (2003): “How Parasites Serve Their Host: A Graphical Analysis of Scalping,” *Public Finance Review*, 31, 694–99.
- Swafford, J.L. (1999): “Arbitrage, Speculation, and Public Policy Toward Ticket Scalping,” *Public Finance Review*, 27(5), 531–40.

- Sweeting, A. (2008): “Equilibrium Price Dynamics in Perishable Goods Markets: The Case of Secondary Markets for Major League Baseball Tickets,” *Mimeo*.
- Thiel, S. (1993): “Two Cheers for Touts,” *Scottish Journal of Political Economy*, 40(4), 447–55.
- Williams, A. (1994): “Do Anti-Ticket Scalping Laws Make a Difference?” *Managerial and Decision Economics*, 15(5), 503–09.
- Yatchew, A. (1997): “An Elementary Estimator of the Partial Linear Model,” *Economics Letters*, 57, 135–43.

Table 1: Summary statistics: Events ($N = 56$)

	Mean	Std. Dev.	Min	Percentiles			
				.25	.50	.75	
Primary Market:							
Tickets sold	18286.20	6831.47	3169.00	13859.00	16920.00	21763.00	34844.00
Tickets comped	184.39	147.12	0.00	60.00	145.00	316.00	494.00
Revenue (000)	1481.14	508.16	266.33	1119.63	1377.48	1912.43	2323.90
Venue capacity	18544.54	6824.16	3171.00	14085.00	17483.00	22087.00	35062.00
Capacity util.	0.99	0.02	0.83	1.00	1.00	1.00	1.00
Average price	90.54	44.35	43.38	54.48	68.21	144.15	187.24
Maximum price	150.13	112.05	47.50	66.65	85.85	307.40	315.50
# price levels	2.71	1.07	1.00	2.00	2.00	4.00	4.00
% first week	0.70	0.14	0.28	0.62	0.73	0.80	0.96
Secondary Market:							
Tickets resold	916.39	543.49	377.00	580.00	704.00	1101.00	3130.00
Resale revenue	103.76	54.18	42.33	65.40	87.48	121.53	295.32
Percent resold	0.05	0.03	0.03	0.03	0.04	0.06	0.17
Percent revenue	0.08	0.06	0.03	0.05	0.06	0.09	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 ($N = 51, 318$)

	Mean	Std. Dev.	Min	Percentiles			
				.25	.50	.75	Max
Resale price	113.23	80.91	3.50	66.25	91.50	135.00	2000.00
Markup	22.80	68.64	-308.65	-0.85	20.60	44.50	1686.40
% Markup	0.41	0.75	-0.98	-0.01	0.32	0.67	14.86
Seat quality	0.61	0.27	0.00	0.37	0.65	0.85	1.00
Days to event	43.45	42.76	0.00	7.00	26.00	76.00	208.00
Sold by broker	0.54	0.50	0.00	0.00	1.00	1.00	1.00
Sold below face value:							
by broker	0.21	0.41	0.00	0.00	0.00	0.00	1.00
by non-broker	0.31	0.46	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: Estimated parameters

Parameter	Notation	Estimate	Std. Error
Consumers' transaction cost	τ^c	61.208	1.417
Brokers' transaction cost	τ^b	5.389	2.032
Curvature	ϕ	1.011	0.166
Mean of willingness to pay	μ_ω	-0.345	0.299
SD of willingness to pay	σ_ω	1.085	0.184
Mean of arrival cost	μ_θ	1.739	0.180
SD of arrival cost	σ_θ	1.554	0.121
Correlation of WTP, arrival cost	$\sigma_{\omega\theta}/\sigma_\omega\sigma_\theta$	-0.522	0.036
Scale for brokers' arrival costs	δ_b	0.169	0.206
Prob(conflict)	ψ	0.014	0.003
Prob(broker)	β	0.004	0.003
Scale factor for event effects	$\bar{\alpha}$	1.043	0.069
Parameter of arrival cost function	c_0	1.061	0.108
Number of bidders in resale auctions	μ_L	2.364	0.488
SD of event-specific demand shock	σ_Δ	0.445	0.132

Table 4: Counterfactual simulations: no resale vs. frictionless resale

	Endogenous Arrival			Random Arrival	
	No resale	Base case	Frictionless	No resale	Frictionless
Gross surplus of attendees	100.0	101.0	111.7	68.5	108.5
Transactions costs incurred	0.0	0.5	0.0	0.0	0.0
Arrival costs incurred	6.9	7.0	8.7	0.0	0.0
Net surplus	93.1	93.6	103.0	68.5	108.5
Primary market revenues	43.1	44.0	45.4	44.1	46.2
Resellers' profits:					
Brokers	0.0	0.1	0.2	0.0	0.3
Non-brokers	-3.6	-3.6	13.0	0.0	10.8
Attendees' net surplus	53.6	53.1	44.4	24.4	25.9

Numbers represent averages across events, with 100 model simulations for each event. Numbers are normalized so that attendees' gross surplus equals 100 in the "no resale" case with endogenous arrival. For the "base case" column, the model is simulated at the estimated parameters. The "no resale" column reflects outcomes when transactions costs are set arbitrarily high; the "frictionless" case reflects outcomes when transactions costs are set to zero and the number of bidders in the secondary market auctions is set to 600. The "endogenous arrival" columns correspond to the model we estimate, in which buyers make strategic arrival decisions in the primary market. In the "random allocation" columns, we simply assign the buyer sequence randomly (and independently of buyers' willingness to pay).

Table 5: Counterfactual simulations: the impact of brokers

	No brokers ($\beta = 0$)	Base case ($\beta = 0.004$)	More brokers ($\beta = 0.02$)	More brokers ($\beta = 0.10$)
Gross surplus of attendees	99.3	100.0	104.8	107.2
Transactions costs incurred	0.4	0.5	0.5	1.0
Arrival costs incurred	6.9	6.9	6.9	6.2
Net surplus	91.9	92.6	97.3	100.1
Primary market revenues	42.7	43.4	47.0	51.7
Resellers' profits:				
Brokers	0.0	0.1	0.7	1.8
Non-brokers	-3.6	-3.6	-3.3	-2.3
Attendees' net surplus	52.8	52.7	53.1	49.7

Numbers represent averages across events, with 100 model simulations for each event. Numbers are normalized so that attendees' gross surplus equals 100 in the "no brokers" case.

Figure 1: Two sample events

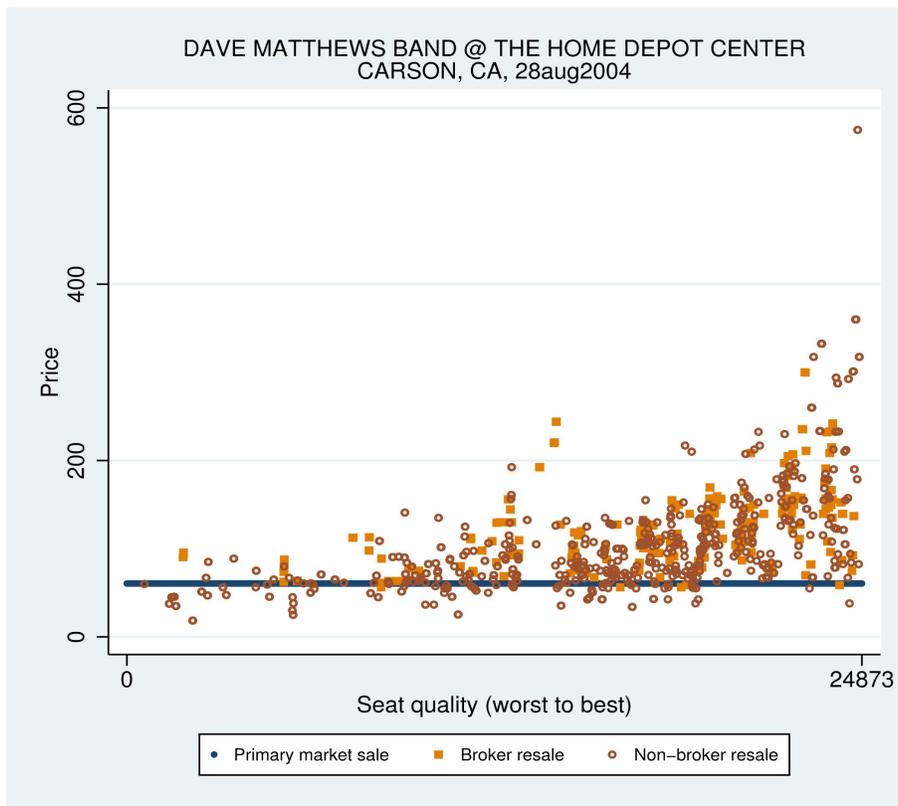
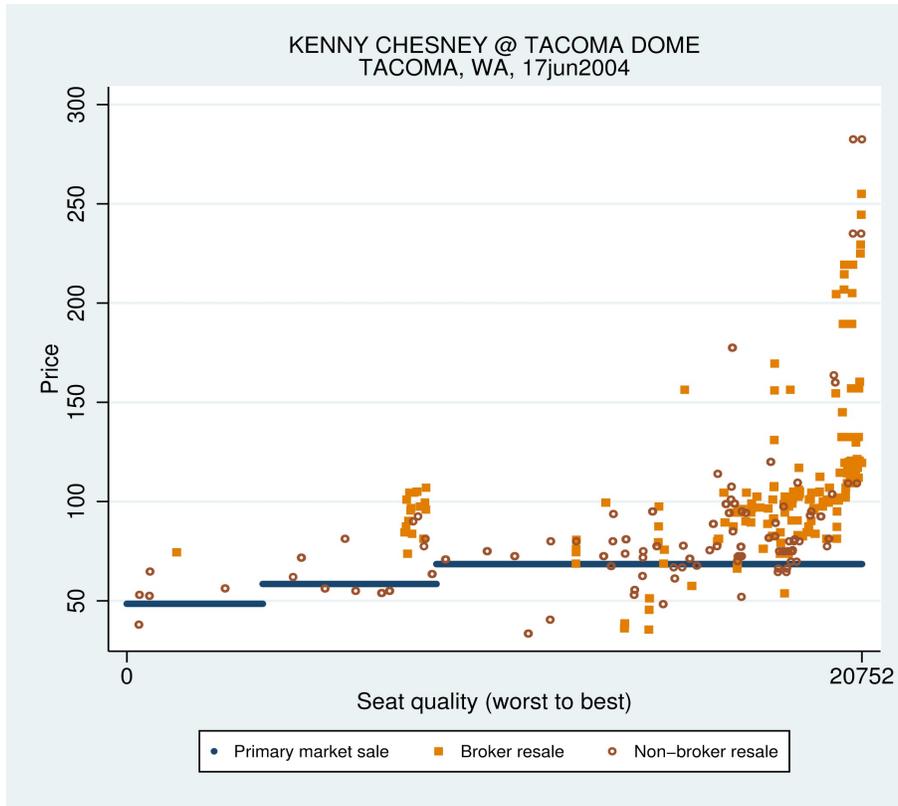


Figure 2: Probability of resale and relative seat quality

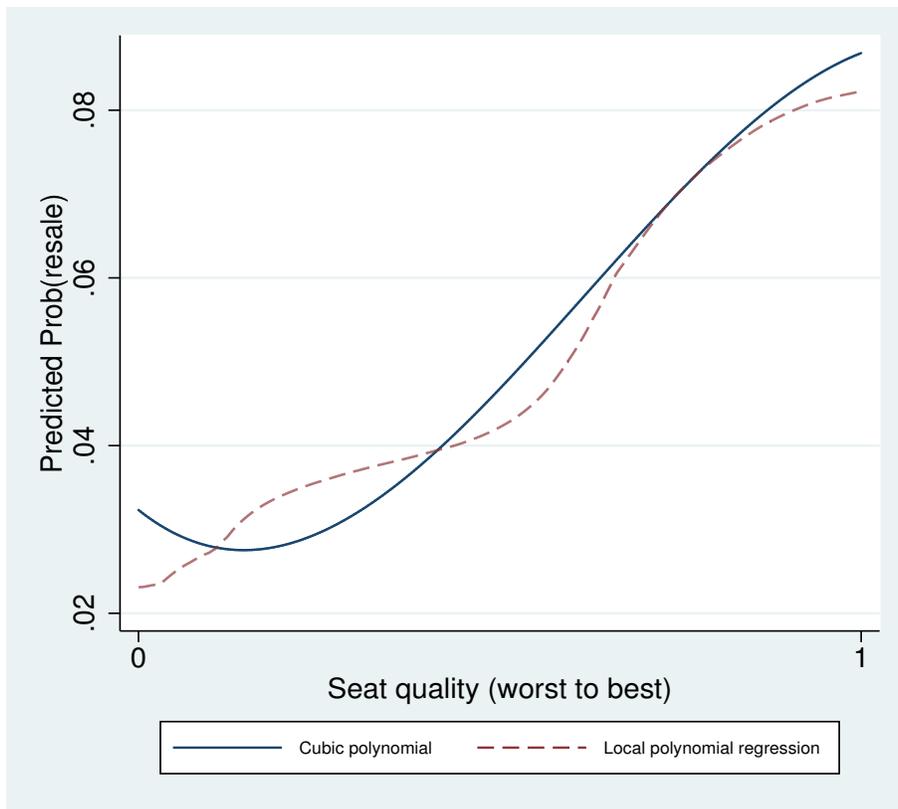


Figure 3: Resale prices and relative seat quality

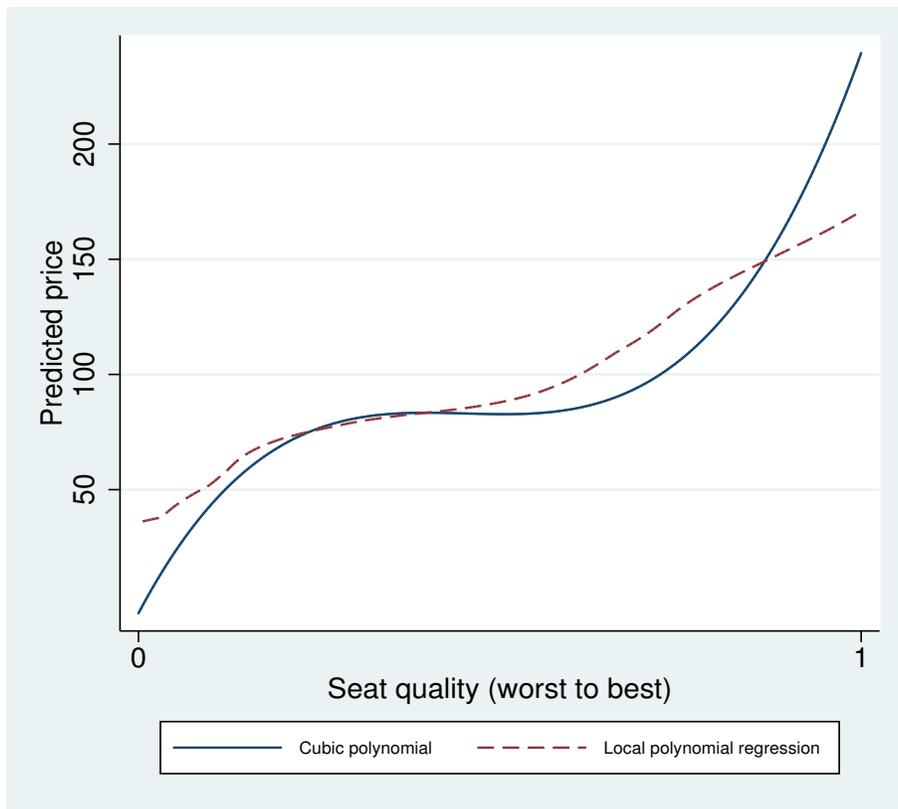
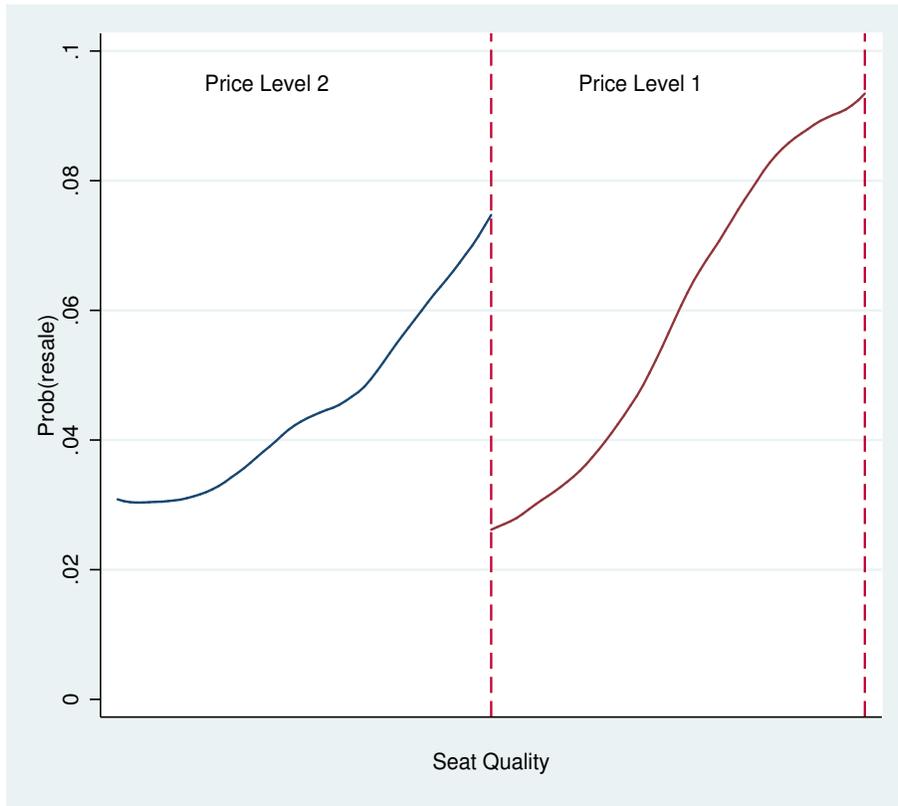
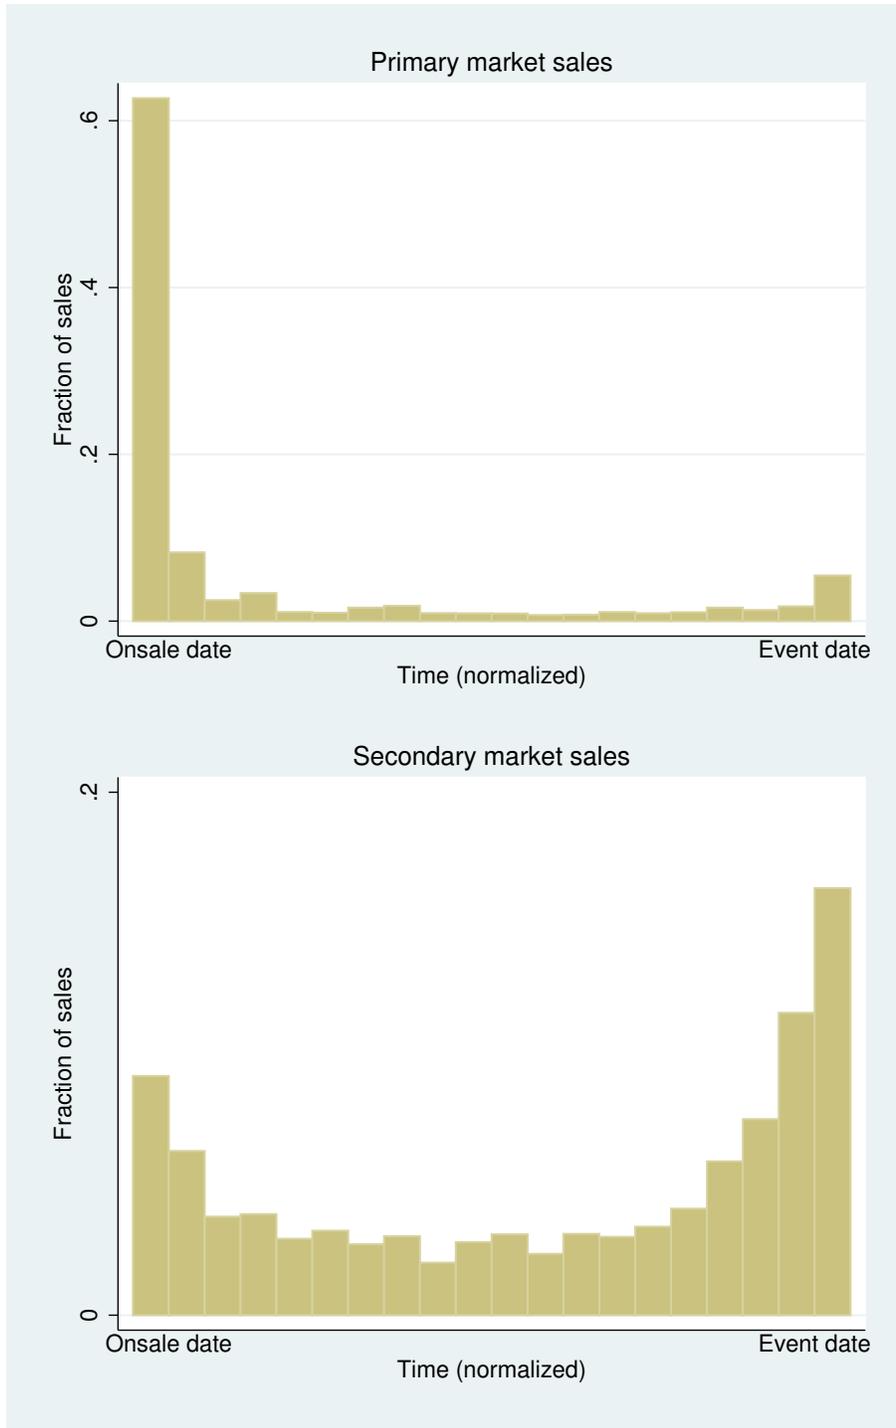


Figure 4: Probability of resale by price level



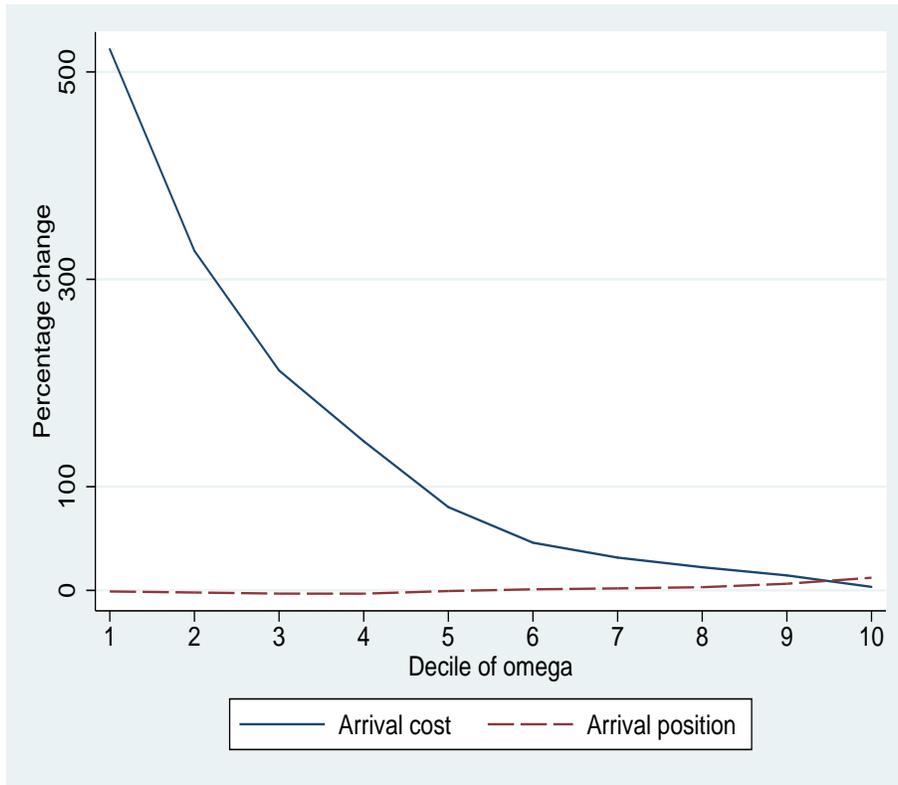
In generating this figure, only events with two or more price levels were used. Relative seat qualities are calculated *within* price level for this figure, and the probability of resale is estimated using kernel-weighted local polynomial regression. So, for example, the probability of resale is on average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

Figure 5: Timing of sales in primary and secondary markets



Time is normalized to make it comparable across events; it is measured as (days since onsale)/(total days between onsale and event). The histogram in the top panel represents the 1,034,353 tickets sold by Ticketmaster; the bottom panel represents the 51,318 tickets resold on eBay or StubHub.

Figure 7: The impact of resale on arrival costs, by willingness-to-pay (ω)



Percentage changes in average arrival cost incurred (solid line) and average arrival position (dashed line), by decile of willingness-to-pay (ω), when we move from a world without resale to a world with frictionless resale. For example, the average arrival costs of buyers with ω 's in the lowest decile increase by over 500% under frictionless resale, but their average arrival position is essentially unaffected.