Primary-Market Auctions for Event Tickets: Eliminating the Rents of "Bob the Broker"*

Aditya Bhave† and Eric Budish‡

September 3, 2012

Abstract

Economists have long been puzzled by event-ticket underpricing: underpricing reduces revenue for the performer, and encourages socially wasteful rent-seeking by ticket brokers. Why not use an auction to set price correctly? This paper studies the recent introduction of auctions into this market by Ticketmaster. By combining primary-market data from Ticketmaster with secondary-market resale value data from eBay, we show that Ticketmaster's auctions work well overall: the auctions substantially improve price discovery, nearly double performer revenues, and, on average, nearly eliminate the arbitrage profits associated with underpriced tickets. We then investigate the importance of bidding experience. The most experienced bidders earn modest profits on average, while inexperienced bidders disproportionately make large bidding mistakes, and then disproportionately never return to bid in another auction. These results suggest that modifying the TM auction design to reduce its strategic complexity would enhance the long-run health of the marketplace.

Preliminary and incomplete. Please do not circulate.

*We are grateful to Kip Levin and especially Debbie Hsu of Ticketmaster, for providing the Ticketmaster auction data and teaching us a great deal about event ticket markets. We are grateful to Adam Juda and Aaron Roth for the Perl scripts we used to obtain eBay data. We thank Susan Athey, Matt Gentzkow, Philip Haile, Ali Hortacsu, Alan Krueger, Julie Mortimer, Roger Myerson, Chris Nosko, Matt Notowidigdo, Ariel Pakes, Canice Prendergast, Phil Reny, Al Roth, Michael Schwarz, Jesse Shapiro, Alan Sorensen, Lars Stole, Chad Syverson, Steve Tadelis, and seminar participants at the University of Chicago and eBay Research for helpful discussions. We thank Sam Hwang, Meru Bhanot and Mika Chance for excellent research assistance. Disclosure: Ticketmaster's primary-market auction data were provided under a confidentiality agreement, without restrictions on the right to publish other than that artists' names be kept confidential in the empirical analysis. Financial support from the University of Chicago Booth School of Business is gratefully acknowledged.

†University of Chicago, abhave@uchicago.edu

‡University of Chicago Booth School of Business, eric.budish@chicagobooth.edu.

"It is nevertheless true that gangs of hardened ticket speculators exist and carry on their atrocious trade with perfect shamelessness." New York Times Editorial [1876].

"Several decades ago I asked my class at Columbia to write a report on why successful Broadway theaters do not raise prices much; instead, they ration scarce seats, especially through delays in seeing a play. I did not get any satisfactory answers, and along with many others, I have continued to be puzzled by such pricing behavior." Gary Becker [1991]

"We're in an industry that prices its product worse than anybody else." Terry Barnes, former chairman of Ticketmaster, in the Wall Street Journal [2006].

1 Introduction

In early 1868, Charles Dickens read from *A Christmas Carol* at the Steinway Hall in New York City. Tickets sold out in half a day at their face value of $2, and reportedly had a secondary-market value of as high as $20; another report indicated that a young boy was paid $30 in gold for a good spot in line [New York Times, November 1867 and December 1867]. This phenomenon of event-ticket underpricing – in which an artist, intentionally or not, sets a price for her event at a level at which demand substantially exceeds supply – predates even Dickens [Segrave, 2007] and is widespread in the present day [Leslie and Sorensen, 2009]. For instance, when the Disney star Miley Cyrus (aka Hannah Montana) first toured the US in 2007-2008, tickets with a face value of at most $64 sold out in approximately twelve minutes, and were then immediately posted on secondary-market venues such as eBay and StubHub at prices that in some instances exceeded $2000 [Levitt, 2008].

Economists have long found this phenomenon puzzling [Becker, 1991, Landsburg, 1993, Rosen and Rosenfield, 1997, Courty, 2003a, Leslie and Sorensen, 2009, Sweeting, 2012]. First, underpricing reduces revenues for the artist. Second, underpricing encourages socially wasteful rent-seeking by ticket speculators.[1] Modern information technology has only exacerbated the scale, and publicness,

[1]It is worth emphasizing that the puzzle is the *combination* of low prices and an active secondary market. Artists may genuinely wish to sell tickets to their fans at a low price, either out of altruism [cf. Che et al., 2011] or to sell complementary goods [cf. Mortimer et al., 2012], but it is hard to argue that artists genuinely wish for ticket speculators to make arbitrage profits. There have been some efforts by artists to set low prices but also make tickets non-transferable, e.g. Miley Cyrus's second tour in 2009 [Waddell, 2009] or the 2012 Summer Olympics in London [Economist, 2012]. There have also been active lobbying efforts to prevent artists from making tickets non-transferable, e.g. the Fan Freedom Project backed in part by eBay and StubHub [Lipka, 2011].

of the rent-seeking.[2] Recent industry estimates are that 20% of all tickets purchased in the primary market are now resold in the secondary market, constituting on the order of $4bn of volume annually. In extreme cases, speculators amass as many as 90% of the tickets available for a particular event.[3]

A natural question is: *why not use an auction?* The basic function of an auction is to discover market-clearing prices for hard-to-price goods. Relative to fixed prices, an appropriately designed auction should improve price discovery, increase revenue, and reduce or eliminate the arbitrage profits between the primary market and the secondary market. This line of thinking led Ticketmaster ("TM"), the largest ticket distributor in the United States, to introduce primary-market auctions as an option for its clients in 2003. TM's Chief Executive Officer remarked at the time:

> "The tickets are worth what they're worth. If somebody wants to charge $50 for a ticket, but it's actually worth $1000 on eBay, the ticket's worth $1000. I think more and more our clients – the promoters, the clients in the buildings and the bands themselves – are saying to themselves *'Maybe that money should be coming to me instead of Bob the Broker.'*" [Nelson, 2003. Emphasis added]

This paper shows that TM's primary-market auctions worked: the auctions substantially improve price discovery, more than doubled artist revenues, and, on average, nearly eliminated the arbitrage profits of "Bob the Broker".

Our empirical research design is very simple. We have proprietary data from TM which indicate the price at which each particular ticket is sold in TM's primary market auction, for all TM concert auctions conducted in 2007 (we will describe the auction rules in detail below, which are interesting

[2]Historically, speculation was a localized activity without significant economies of scale. Amassing tickets in the primary market required either "diggers" – groups that "push and intimidate their way to the front of the line..." – or "ice" – bribes paid to a venue operator or box-office employee. The secondary market involved "scalpers", individuals who stand outside an event buying and selling tickets [New York Attorney General, 1999]. In the present day, speculation does not require a local presence and has substantial scale economies: tickets can be amassed in the primary market using "software bots" and sold on the secondary market on websites such as eBay and StubHub [Zetter, 2010; Ticketmaster Corporate Blog, June 2011]. On the speed and publicness of modern-day speculator rent seeking, the following quote from the Arkansas Attorney General's office is instructive: "All hell broke loose with Hannah Montana. The tickets were gone in 12 minutes and when people turned around, they were selling at online sites for as much as 10 times the face value." [Rosen, 2007]

[3]The 20% figure is from a blog post by Ticketmaster's CEO on 08/12/2011 [Ticketmaster Corporate Blog, August 2011], and the $4bn figure is from both [Mulpuru, 2008] and a Ticketmaster estimate. The 90% figure is a Ticketmaster estimate, based on software used to detect the "software bots" mentioned in the previous footnote. Some press reports cite volume estimates as high as $10bn annually [Stecklow, 2006]. To give a sense of the growth of secondary market activity, Leslie and Sorensen [2009] find that the rate of resale was on the order of 5% in summer 2004. TM also writes in its corporate blog [June 2011] that the use of "bots ... [that] cut in line ahead of you and scoop up large quantities of tickets ... is on the rise."

in their own right). The data cover 23 distinct concert tours and 579 distinct concerts. We also have data from TM which indicate, for each ticket sold by auction, what the fixed price would have been if not for the auction; this is possible since only a relatively small number of tickets per event were sold by auction, so it was always the case in our data that some tickets in the same pricing tier as those sold by auction were sold by fixed price.[4] Last, we have data from eBay on the secondary-market values of tickets sold by TM's auction. Specifically, we scraped all instances where a ticket substantially identical to a ticket sold in TM's auction – a ticket to the same event, on the same date, in the same section and row of the venue – subsequently sold on eBay.[5] Given these three types of data it is straightforward to calculate the TM auctions' effect on price discovery, revenues and arbitrage profits.

Our main empirical finding is that the mean arbitrage profits in our matched sample – i.e., the mean difference between the eBay secondary-market price (net of fees) and the TM primary-market price of substantially identical tickets – is \$4.27, or 1.5% of the average TM primary-market auction price of \$277.49. The 95% confidence interval of this estimate is [-\$10.31, \$16.71] if we cluster standard errors at the concert level, or [\$1.11, \$7.57] if we do not cluster. Thus, mean arbitrage profits are economically indistinguishable from zero, and, depending on the specification, statistically indistinguishable as well. Relative to the (actual, non-optimal) face values they replaced, auction prices are substantially more informative of secondary market values, with the R^2 of the regression of secondary-market values on primary-market prices increasing from 0.235 to 0.647, and nearly doubled performer revenue, from \$8.8mm to \$16.9mm.[6] Put differently, the auctions not only discover significantly different prices from the counterfactual face values, but furthermore these prices are essentially correct on average.

Figure 1 displays the distribution of this difference between eBay secondary-market prices and TM primary-market auction prices. Note that the mean of the distribution is close to zero, as described above. There are two other interesting features of the distribution to highlight. First, there is substantial variance: while the mean arbitrage profits are close to zero, there are of course

[4]For example, for many events all "floor" seats have the same fixed price (i.e., are in the same pricing tier), while only a fraction of floor seats are sold by auction.

[5]The reason that we match tickets at the level of section and row, rather than the precise seat within the row, is that for privacy reasons eBay listings do not indicate the specific seat. This is standard practice in the secondary market. See Section 4.3 for more details.

[6]The large magnitude of the revenue difference is consistent with Leslie and Sorensen [2009]'s finding that high-quality tickets, which are the focus of TM's auctions, is where there is the most underpricing.

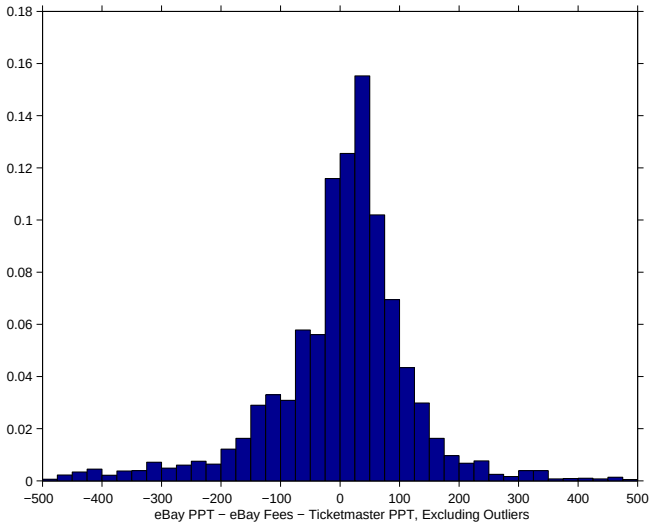


Figure 1: Distribution of "Arbitrage Profits", calculated as the difference between the eBay secondary-market price per ticket (PPT) and the Ticketmaster primary-market price per ticket, net of eBay transaction fees. For more details see Section 5.

specific tickets where the secondary-market value turns out to be substantially higher than the primary-market auction price, and vice versa. Second, the distribution is slightly asymmetric: Specifically, the modal outcome of small positive profits ($25-$50 per ticket) is greater than the mean outcome of essentially zero profits, and the left tail (large losses) has more density than the right tail (large gains). We find that left-tail outcomes are disproportionately likely to correspond to tickets purchased by inexperienced TM auction participants, whereas the 1% most experienced bidders, who account for 16% of total volume, account for a disproportionate share of the positive-profit outcomes and earn small positive profits on average, of $18 per ticket. The $18 per ticket can be interpreted as a return for the time, effort, and risk associated with ticket speculation (cf. Courty [2003a,b]). While we cannot say whether $18 per ticket represents an excess rent, we emphasize that it is an order of magnitude smaller than the $135 mean rent associated with purchasing these same tickets at their counterfactual face values. We also find that left-tail outcomes are disproportionately likely to correspond to eBay sales that occur in the last few weeks before an event occurs, whereas resale profits are higher in eBay sales that occur in the first few weeks after TM's auction. This finding is consistent with the declining price phenomenon documented by Sweeting [2012].

So far we have left vague the specific details of TM's auction design. It turns out that the auction

TM designed is a variant on the position auctions that Google and other search engine firms use for keyword advertising [Edelman et al., 2007, Varian, 2007]. This similarity makes sense, because in each case the auction is for a set of vertically-differentiated goods: in the keyword advertising case, the goods are advertising slots of varying proximity to the top of the search page (e.g., first slot, second slot, etc). In TM's case, the goods are tickets of varying proximity to the stage (e.g., first row, second row, etc.). There are two main differences between TM's auction and the most well known and widely used position auction format, the Generalized Second Price (GSP) auction. The first difference is that in TM's auction successful bidders pay their own bid, rather than the next highest bid; that is the auction is more like Generalized *First* Price (GFP) than GSP. TM perceived that the benefit of using pay-your-own-bid is that it is simpler to explain to its customers. The second difference is the nature of bids: in TM's auction, bids are a per-ticket dollar amount, whereas typically in GSP bids are a per-click dollar amount. This difference also makes sense for simplicity reasons, because it would be difficult to bid on a "per-unit-quality" basis, whereas in the keyword advertising setting it is natural to bid on a per-click basis.

While the TM auction design is simple to explain it is complex strategically, and we find some evidence in the data that bidders make occasional large mistakes. These mistakes are rarely made by experienced bidders, and are disproportionately made by first-time bidders, who then disproportionately never return to bid in another auction. Thus, strategic complexity has a detectable negative effect on the long-run health of the TM auction marketplace. Additionally, strategic complexity might have a negative effect on the health of the marketplace which we cannot detect in our data, namely that some potential bidders opt never to participate in the first place. We discuss a simple modification to the TM auction design that might mitigate this complexity, which is analogous to moving from GFP to GSP in the position auctions context (cf. Edelman and Ostrovsky [2007]). We also discuss a more ambitious modification that yields a strategyproof auction design.

The remainder of this paper is organized as follows. Section 2 provides institutional background. Section 3 describes Ticketmaster's auction design. Section 4 describes our data. Section 5 presents our main empirical results. Section 6 discusses potential improvements to the TM auction design. Section 7 concludes.

2 Institutional Background

"Primary market" refers to the original sale of tickets to an event, by or on behalf of the event organizer. Ticketmaster, established in 1976, is the world's largest primary-market ticket distribution company. In 2009, Ticketmaster sold more than 130 million event tickets valued at $8.0 billion, on behalf of clients including venues, promoters, sports leagues and teams, and museums and cultural institutions. Tickets are typically sold at fixed prices that vary coarsely with seat quality, and typically go on sale months in advance of an event [Ticketmaster 2010 Form 10-K].

"Secondary market" refers to the resale of tickets purchased in the primary market. Ticketmaster recently estimated that 20% of all tickets purchased from Ticketmaster in the primary market are subsequently resold on the secondary market [Ticketmaster Corporate Blog, June 2011]. The e-commerce research firm Forrester Research has estimated secondary market dollar volume as on the order of $4bn annually, and other estimates range as high as $10bn annually [Mulpuru, 2008, Stecklow, 2006]. The largest forums for secondary market activity are eBay and StubHub; Ticketmaster itself entered the secondary market in 2002 with its launch of TicketExchange, and then increased its presence in 2008 with its purchase of TicketsNow [Ticketmaster 2010 Form 10-K]. The secondary market is most active in the final days before the event occurs [Sweeting, 2012].

An important recent study by Leslie and Sorensen [2009] examines the welfare effects of the secondary market, and finds empirical evidence of substantial costs and benefits and resale. The main benefit is that it enables Pareto-improving reallocation of tickets, e.g., resale by fans who no longer can attend the event. The main cost is the rent-seeking activity that the possibility of resale encourages in the primary market. In Leslie and Sorensen [2009]'s analysis, if price could be set correctly in the primary market such that rent-seeking activity is eliminated, the main cost of allowing resale would be eliminated as well. Our paper suggests that this is possible, via auctions.

While secondary-market activity has been a part of the event-ticket market for a long time (see the quotes in the introduction), its scale seems to have increased dramatically with the rise of the internet.[7] There are at least three reasons. First, the internet has lowered the costs of amassing tickets in the primary market.[8] Second, the internet has lowered the cost of reselling tickets in the

[7] See the New York Attorney General's 1999 report "Why Can't I Get Tickets?" [New York Attorney General, 1999], for an excellent primer on pre-internet ticket reseller tactics.

[8] For instance, Ticketmaster writes on its corporate blog: "There continue to be nefarious online scalpers who use sophisticated tools – often known as bots – to cut in line ahead of you and scoop up large quantities of tickets, only

secondary market. Third, the internet has made it easier to skirt state rules on ticket reselling (cf. Courty [2000, 2003a], Connolly and Krueger [2006]).

Technology has also changed the publicness of the secondary market, e.g., any ordinary fan can now look up the secondary market value of their tickets on eBay. Roth [2007] speculates that this may have caused a decline in the "repugnance" associated with charging high prices for tickets in the primary market, a trend that has manifested both in the use of auctions and in the use of higher fixed prices than in previous eras (cf. Connolly and Krueger [2006]).

3 Ticketmaster's Primary-Market Auction

In 2003 Ticketmaster introduced auctions as a primary-market pricing method, alongside fixed price. In this Section we describe the rules of the auction in detail and briefly discuss its formal theoretical properties.

Which Tickets are Auctioned? For any particular event, the determination of whether and which tickets to sell by auction, and which to sell by fixed price, is made by TM's client. In our data, an average of about 75 tickets are auctioned per concert, with a max of 862 tickets. The auctioned tickets are always of high quality, often in the first few rows of the venue, allowing the auction to be positioned in TM's marketing efforts as "premium seat auctions".[9] This decision to focus on high-quality tickets is consistent with Leslie and Sorensen [2009]'s finding that high-quality tickets is where there is the most underpricing and inefficient rent-seeking. TM and the client organize the auction tickets into discrete quality groups, typically by rows. For instance, in the auction depicted in Figure 2, the first quality group is "Section A3, A4 or A5, Row 2", the second "Section A3, A4 or A5, Row 3", etc. TM and the client also rank the tickets by quality within each group. For instance, within Row 2, tickets in Section A4 are ranked above those in Sections A3 and A5 because Section A4 is more centrally located. The groups are designed however so that quality heterogeneity within a group is small.

to turn around and sell them to fans at many times the face value of the tickets. The use of these bots is illegal, it violates our terms of use, and it is on the rise. Worst of all, these bots prevent you from getting a fair shot at tickets to the event you want to see live" [Ticketmaster Corporate Blog, June 2011].

[9] One issue with expanding the auctions to include a wider variety of seats is that this would require an auction design with a more expressive bidding language. See Section 6.1 for discussion.

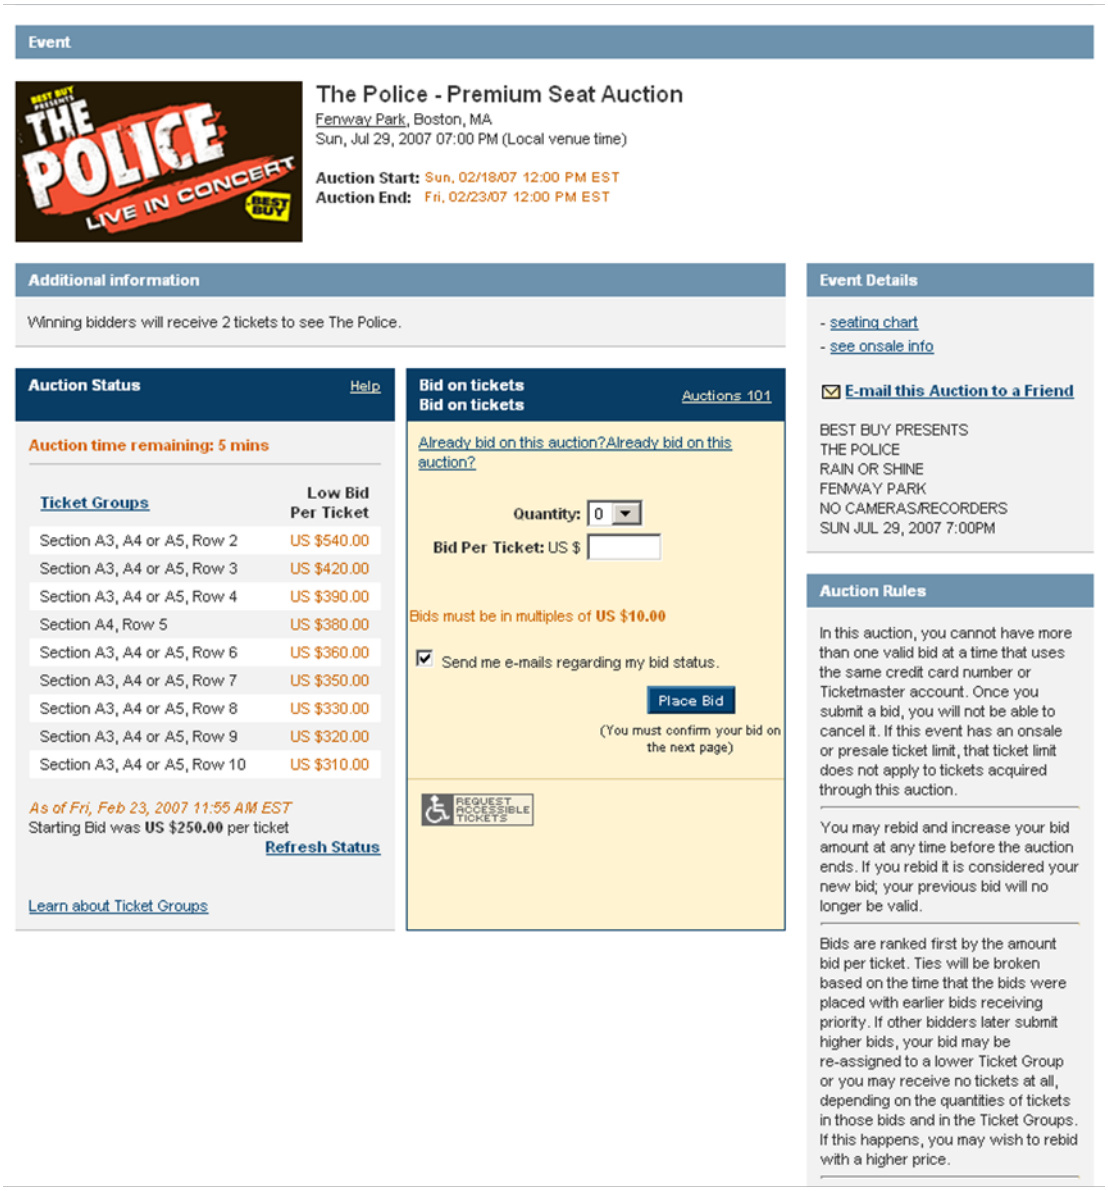


Figure 2: User interface for Ticketmaster auction for a concert by The Police in Fenway Park, Boston, July 29, 2007.

Auction Rules The auction itself lasts for several days, starting and ending at a pre-announced fixed time. The auction dates typically are timed to coincide with the sale of other tickets by fixed price, for marketing reasons; e.g., in the auction depicted in Figure 2, the auction opens on a Sunday and ends on a Friday, while the bulk of fixed-price tickets went on sale on the intervening Tuesday.

The auction has a non-zero per-ticket Starting Bid, which is set to be approximately equal to the fixed-price face value of similar quality tickets.[10] Bids consist of a per-ticket dollar amount and a desired number of tickets, e.g. 2 or 4. Bidders can increase their bid amount at any time throughout the auction, but bidders are not allowed to lower or retract their bids. At the conclusion of the auction, bids are sorted in descending order, with the highest bid winning the best tickets within the highest quality group, the next highest bid winning the next best tickets in the highest quality group, etc. Ties are broken by order of bid receipt. Successful bidders pay their bid amount; losing bidders pay zero.

Over the course of the auction, bidders can view the current market-clearing prices by group. For instance, in the auction depicted in Figure 2, there are enough bids of $540 and higher to fill the first quality group, enough bids of $420 and higher to fill the first two quality groups, etc. Additionally, bidders receive email notifications whenever their current bid's tentative assignment drops down a quality group. For instance, if the cutoff for the first quality group just increased from $530 to $540, any bidders of $530 will just have received a notification.

Relationship to Position Auctions Ticketmaster's auction design is closely related to the position auctions widely used by internet search engines to allocate advertising positions related to a particular search keyword (cf. Edelman et al. [2007], Varian [2007]). There are two central auction design differences between TM's auction and the most well known and widely used position auction format, the Generalized Second Price (GSP). First, in TM's auction successful bidders pay their own bid, rather than the next highest bid. Second, bids are a per-ticket dollar amount, whereas typically in sponsored search auctions bids are a per-click dollar amount.

In Appendix B we modify the position auctions model of Edelman et al. [2007] and Varian [2007]

[10] Auction bids are inclusive of convenience fees, whereas fixed-price tickets have a face value exclusive of convenience fees. The Starting Bid is typically set equal to the face value of similar quality tickets, plus convenience fees, rounded to a multiple of $10 or $25. For instance, in the auction depicted in Figure 2, the face value for tickets in the same sections as the auction was $225, convenience fees were about $22, and the auction Starting Bid was set to be $250.

to account for these two differences, and derive two basic results.[11] First, the TM auction has a monotonic Bayes-Nash equilibrium, which implies (via Myerson's lemma) that the TM auction is revenue equivalent to GSP and that, if the reserve price does not bind, the allocation is efficient. Second, when we add free entry by speculators to the model we find that entry eliminates arbitrage profits for speculators. We interpret these results as formalizing that TM's auction design is sensible, though it does have some weaknesses that we will turn to below in Section 6. The results also can be intepreted as providing general support for the use of auctions over fixed prices, since fixed-price mechanisms satisfy none of these properties.

2007 Rules Change At the time bidders submit their bid, they are uncertain what quality tickets, if any, they will win. For instance, a bidder who bids $600 in the auction depicted in Figure 2 would win tickets in the first ticket group if the auction closed immediately, but might win tickets in any of the other ten ticket groups, or none at all, depending on what other bids are submitted subsequently to his. In principle, if a bidder wins tickets considerably worse than expected, the bidder could end up with negative realized surplus. In our stylized model (cf. Appendix B), this phenomenon can occur in equilibrium.

In mid-2007 TM modified its auction rules in response to concerns about this issue.[12] Under the modified rules, bids consist of a dollar amount and a minimum acceptable quality group. For instance, a bidder could submit the bid $600 and indicate that it is only valid for the first two quality groups. As before, bidders may increase their bids at any time, and also they may decrease their quality threshold at any time. At the end of the auction, as before, bids are sorted in descending order, with the highest bid winning the highest-quality tickets, etc. The only modification is that if a bid is reached where the quality that would be awarded is below the bidder's threshold, then that bid is skipped.

We will discuss the rules change further in Section 6.1.

[11] As in Edelman et al. [2007], we consider both sealed-bid and ascending versions of the model.

[12] One practical issue with using an auction that leaves some bidders with negative surplus is that disappointed bidders can dispute the credit card charge associated with their "successful" bid. Credit card chargebacks are costly and time consuming.

4 Data

Our data come from two sources: primary-market auction data provided by Ticketmaster, and secondary-market resale value data scraped from eBay. Sections 4.1 and 4.2 describe each dataset in turn. Section 4.3 describes how the datasets are matched.

4.1 Primary-Market Auction Data

Our primary-market auction data cover all Ticketmaster auctions for concert tours that started in 2007. There are 23 concert tours, 581 concerts, and 764 auctions.[13][14] The concerts took place between March 2007 and April 2008, and the corresponding auctions were conducted between January 2007 and December 2007.

Although the original dataset includes all bids, our analysis will focus on winning bids (22,366 in total). For each winning bid, we observe the following bid-level variables: customer identification number, bid amount, number of tickets (2 or 4), time of bid, and the section, row and seat numbers assigned to the bid. We also observe the following auction-level variables: artist, event date, event city, starting bid, ticket face value, starting time, ending time, and whether the auction was conducted under the new or old rules. We caution that the ticket face value should not be interpreted as the optimal fixed price, but rather as the actual price set by the artist for tickets in the same section as the auctioned tickets. For instance, in the auction depicted in Figure 2, the face value for all tickets in the relevant sections of Fenway Park was $246.60.

4.2 Secondary-Market Resale Value Data

During the time period from January 2007 to April 2008 we used Perl scripts, one for each of the 23 concert tours covered in our primary-market data from Ticketmaster, to obtain all completed listings from the eBay category "Event Tickets" that included the artist's name in the title. This resulted in a dataset consisting of over 300,000 html files, one for each eBay listing. We then used a

[13] For most concerts there is just a single auction, but in some cases there are 2 or more auctions for the same event. For instance it is somewhat common for there to be a separate auction for tickets in the first row. This can be understood as an auction design response to the large perceived difference in quality between the first and second rows; e.g., under the old rules, such a difference would cause there frequently to be bidders with negative realized surplus.

[14] We drop all concerts in Canada. These concerts comprise less than 2% of our primary market dataset and about 0.1% of our final, matched dataset.



Figure 3: eBay auction webpage for tickets that were sold in the Ticketmaster auction shown in Figure 2

separate Perl script to extract several kinds of data from each html file. First, there are event-level data such as artist, event date and city. Next, there are data on the specific tickets being auctioned, specifically section, row, and number of tickets.[15] Last, there are data on the auction design choices the seller made, as well as the outcome. Design choices include the auction's opening bid and the Buy-it-Now price, if the seller chose to use one.[16] The main outcome variables are whether or not the item sold, and at what price (per ticket).

For most of our analysis we use the eBay price per ticket net of eBay transaction fees. At the time of our data, these fees were roughly 4%.[17]

An example of an eBay auction webpage is depicted as Figure 3. This eBay listing was for a

[15]For both event-level data and ticket-level data, our Perl script exploits the fact that sellers post the information we seek in a structured and consistent way, thanks to what eBay called Category Specific Information at the time of our data, and presently calls Item Specifics.

[16]For additional details on the use of eBay data in economic research see Bajari and Hortaçsu [2003, 2004].

[17]At the time of our data, eBay's fee schedule for Final Value Fees was 5.25% of the first $25, 3.25% of the amount between $25 and $1000, and 1.5% of any amount above $1000. There are also insertion fees, or eBay listing fees, which depend on the reserve price. We ignore PayPal fees, since we cannot observe which eBay transactions occurred via PayPal; at the time of our data, PayPal fees were just under 3%. For a sale of a pair of tickets with a per-ticket sale price of $290, roughly the average in our matched sample, with a reserve price of $200 per ticket, the total fee is $22.95, or almost 4% of the $580 transaction value.

pair of tickets in Section A3, Row 3 to see the Police at Fenway Park on July 29, 2007; this is the same concert whose TM auction webpage we depicted above in Figure 2. This eBay listing resulted in a total sale price of $999.99, or $499.995 per ticket.

Unlike the primary market data, we do not observe seat numbers in the eBay data. For instance, notice in Figure 3 that Section and Row are data fields that eBay allows the seller to fill in (and that this particular seller did fill in), but that there is no such data field for seat number.[18] Thus when comparing the primary and secondary markets, we conduct our analysis at the section-row level, as described in Section 4.3.

We drop all observations in which just a single ticket was sold (as opposed to a pair or more), because prices for individual tickets are not representative of per-ticket prices for groups of tickets.[19] Such observations comprise about 8% of our eBay listings. We also drop observations in which the seller elected to use eBay's "Dutch auction" format for selling variable quantities of tickets. Such observations comprise about 3% of our eBay listings.

4.3 Matching Primary and Secondary Market Data

In this section we describe our procedure for matching the Ticketmaster primary-market data to the eBay secondary-market data. There are three specific issues that are important to highlight.

First, the eBay data indicate the section and row in which the auctioned tickets are located, but not the precise seat numbers. This is standard practice in the secondary market for tickets, both for seller privacy reasons and because quality heterogeneity within a section-row is usually of negligible importance relative to the importance of the section and row information.[20][21] For this reason, we match the two datasets at the level of the concert-section-row ("c-s-r").

[18] Listing tickets at the Section and Row but not Seat level is a common practice on all of the secondary market websites of which we are aware. See further discussion of this issue in Section 4.3.

[19] Intuitively, individual tickets are likely to be less valuable than groups of tickets because concerts are social events that fans attend in pairs or larger groups. To document that individual tickets are systematically less valuable, we regress price-per-ticket (for successful auctions) on a dummy for individual sales, controlling for concert-section-row fixed effects. We find that the average price of individual tickets sold on eBay is 36% lower than the average price-per-ticket in multiple ticket sales.

[20] In addition to typically being of negligible importance, seat data is also typically difficult to interpret. For instance, for the auction depicted in Figure 2, the Fenway Park concert seat map indicates that Section A4 is slightly more centrally located than Sections A3 and A5, and it is obvious that "Row 2" is higher quality than "Row 3", but information on what seat number is most centrally located within Section A4 - Row 2 is not readily available. As it turns out there are 24 seats within this row, and seats 12 and 13 are the most centrally located.

[21] While seat data being of negligible importance is typical, there are a handful of venues where heterogeneity in seat quality within a row is of sufficient importance that Ticketmaster demarcated distinct quality groups within a row in the auction. One of our robustness exercises omits these venues from the analysis; the results move very little.

Second, eBay section and row data is input by eBay sellers, and is non-standardized. For instance, eBay sellers typically use the string "1" in the Row field to describe tickets that are in Row 1, but we also observe entries of "#1", "**1**, "1st", "1 !!!!", "First", "one", "1 WOW!", and dozens of others. We handle this issue as follows. First, we create a dictionary that translates all observed eBay section and row input strings into standardized terms; e.g., all of the Row entries listed above get translated into "1". We then match the eBay data to the Ticketmaster data at the level of concert-section-row, using this dictionary. Last, we hand check all matches and remove any that we can identify as false positives. This process may miss some matches, which would reduce statistical power.

Third, when a particular concert-section-row tuple has multiple TM primary-market auctions and/or multiple eBay secondary-market transactions, an issue arises as to how exactly to match the two sets of transactions. To illustrate, suppose that for a particular c-s-r we observe the transactions depicted in Figure 4. In both the primary- and secondary- markets, the average transaction price is $200, hence average arbitrage profits are zero. However, there is variation in these arbitrage profits; for instance a speculator who bought tickets on Ticketmaster for $150 would have realized positive profits, while a speculator who bought tickets on Ticketmaster for $250 would have realized losses. Our main specification for the analyses in Section 5 performs this match by aggregating eBay transactions at the c-s-r level. Specifically, for each c-s-r, we calculate the mean price over all eBay transactions in the c-s-r, and then match this mean secondary-market value to each TM primary-market transaction. In the example depicted in Figure 4, this approach leaves us with three matched observations: one in which the primary-market price is $150 and the secondary-market value is $200, one in which the primary-market price is $200 and the secondary-market value is $200, and one in which the primary-market price is $250 and the secondary-market value is $200. For robustness, we also consider a specification that instead aggregates the TM transactions at the c-s-r level and treats each eBay observation separately, and a specification that aggregates both TM and eBay transactions at the c-s-r level. The advantage of our main specification is that it allows us to analyze the TM primary-market data at more granularity than the c-s-r level; e.g., we can ask whether experienced bidders obtain better auction outcomes than inexperienced bidders.

Table 1 provides summary statistics on our matched dataset. The Data Appendix provides more details on the matching process.

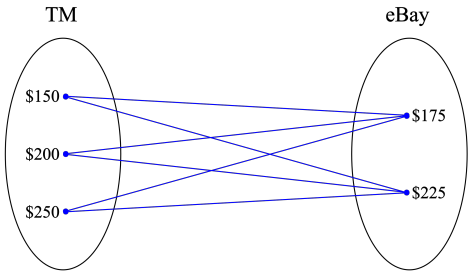
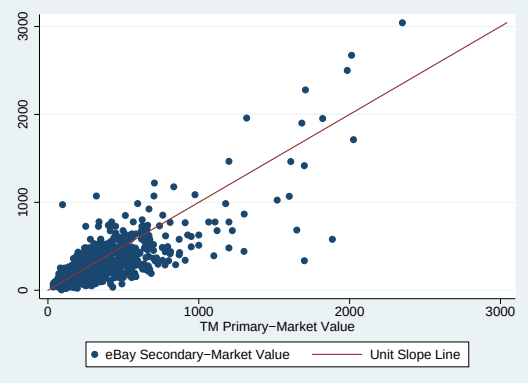


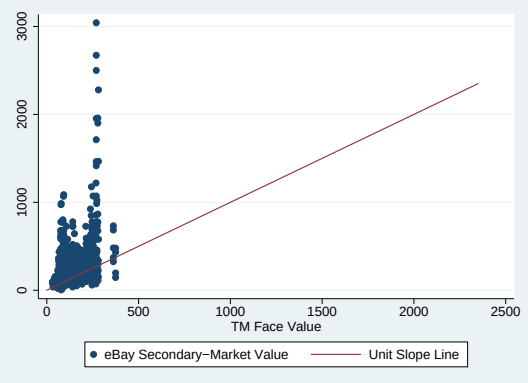
Figure 4: Hypothetical data from the TM primary market and the eBay secondary market.

	Full TM Data Set	Matched Data Set	% Matched	
Concerts	579	453	78.2%	
C-S-R combinations	5,815	1,591	27.4%	
TM transactions	22,366	7,973	35.6%	
eBay transactions	N/A	3,035	N/A	
Matched Data Set	Mean	Std Dev	25th Perc	75th Perc
TM transactions per C-S-R combination	5.01	7.31	2	6
eBay transactions per C-S-R combination	1.91	3.25	1	2

Table 1: Summary statistics for matched primary- and secondary-market data. "C-S-R" stands for Concert-Section-Row.



(a) eBay price per ticket, net of fees, versus TM auction price per ticket



(b) eBay price per ticket, net of fees, versus TM face value

Figure 5: Relationship between aftermarket value and TM prices

5 Do Primary-Market Auctions Discover Secondary-Market Values?

5.1 No Arbitrage

Figure 1 (in the introduction) displays the distribution of the resale profits associated with buying tickets in the Ticketmaster primary-market auction and reselling them in the eBay secondary market, using our matched dataset described in Section 4.3. The mean resale profit is $4.27, or 1.5% of the mean primary-market auction price of $277.49. A 95% confidence interval of this estimate, clustering errors at the concert level, is [-$10.31, +$16.71];[22] unclustered, the confidence interval is [$1.11, $7.57]. Thus, the arbitrage profits associated with buying tickets in the TM primary-market auction are economically small, and, in our preferred specification, statistically indistinguishable from zero. Robustness tests reported in Appendix A suggest that, if anything, the $4.27 estimate is too high and the mean arbitrage profits are slightly negative.[23]

It is instructive to also look directly at the data underlying Figure 1. Figure 5, Panel (a) presents a scatterplot of our matched dataset, with the horizontal axis denoting price per ticket in the TM primary-market auction, and the vertical axis denoting price per ticket in the eBay

[22]The reason to cluster standard errors at the concert level is that it seems natural to think of each concert as its own market. Although tickets for some concerts are sold in multiple auctions, we expect unobservables such as secondary market demand to be correlated within a concert.

[23]Over our three matching specifications, the 95% confidence intervals admit estimates of net arbitrage profits ranging from -26.29 to +16.71. See Appendix A.1.

eBay Secondary-Market Value	(1)	(2)	(3)
TM Face Value	1.75		0.36
	(0.29)		(0.18)
TM Primary-Market Value		0.84	0.79
		(0.07)	(0.09)
Constant	25.23	47.86	10.21
	(32.94)	(18.28)	(18.45)
R^2	0.235	0.647	0.654
Akaike Info Criterion	13.53	12.76	12.73
Schwarz Info Criterion	36,237	30,079	29,924

Table 2: Price informativeness regression results

secondary market. The vertical distance between a point and the 45-degree line represents the profits associated with resale. The reasonably close fit of the data to the 45-degree line conveys both that primary-market auction prices are informative of secondary-market prices ($R^2 = 0.647$), and that average arbitrage profits are small. The regression best fit line is not the 45-degree line though. Rather, the regression best fit features a positive constant ($47.86) and a slope less than one (0.84), with both differences statistically significant (column (2) of Table 2). These regression coefficients suggest that higher-priced auction tickets tend to be systematically overpriced relative to secondary-market value. This is another view of the same phenomena depicted in Figure 1, namely the positive mode and the fat left tail.

To put these results in perspective, consider Figure 5, Panel (b). This figure also presents a scatterplot of the matched dataset, but instead of using TM primary-market auction prices on the horizontal axis uses tickets' face values. Face values should not be interpreted as optimal fixed prices, but rather are the actual prices set by the performer, in consultation with TM, for all tickets in the same section as the auctioned tickets.[24] As can be seen from the figure, ticket face values are substantially less informative of secondary-market prices ($R^2 = .235$), and are systematically underpriced, which manifests in the figure as most of the mass being above the 45-degree line.[25] In aggregate over all of the tickets in our TM data, the TM auction raised $16.9mm of revenues, whereas the face values would have raised just $8.8mm. These results confirm the basic benefits of using auctions over fixed prices: auction prices are more informative, raise more revenue, and

[24] We will discuss the tradeoff between auctions and optimal fixed prices below in Section [6].

[25] Observe that uninformativeness and underpricing are distinct phenomenon. For instance, if face values are always one-half of secondary-market value face values would be highly informative ($R^2 = 1$) despite there being systematic underpricing.

nearly eliminate the arbitrage profits between the primary market and the secondary market.

In the remainder of this Section we explore the no arbitrage finding in more detail on various sub-samples of our data. In particular, we seek to understand what drives the asymmetries in the distribution of arbitrage profits, and the deviation of the regression best fit line from the 45-degree line.

5.2 Experienced vs. Inexperienced Bidders

Our results in Section 5.1 show that primary-market auctions nearly eliminate the average arbitrage opportunity associated with systematically underpriced tickets. If "Bob the Broker" purchases a random ticket in the TM auction, and then resells in the secondary market, he earns negligible profits.

However, experienced bidders may have specialized knowledge about which tickets to purchase. Hence, to fully assess whether the auctions eliminate the rents of Bob the Broker, we need to look separately at the arbitrage profits of experienced and inexperienced bidders. If arbitrage profits are small on average, but large for professional resellers, this would cast the results of the previous section in a different light. It is especially important for us to assess this possibility given the asymmetric distribution of arbitrage profits as depicted in Figure 1, which features a strictly positive mode and a fat left tail. Perhaps it is professional resellers who constitute the modal outcomes, but occasional bidding errors by casual fans create the fat left tail and the appearance of no arbitrage.

We exploit the fact that our TM data contains a unique bidder identifier to define a simple measure of experience, namely, the number of distinct auctions that the bidder has won. We define a bidder as experienced if the bidder wins at least 10 TM auctions (overall, not just restricted to the matched data). Such bidders account for 1% of the bidders who appear in our matched data and roughly 16% of the transaction volume. Bidders who win between 1-9 auctions we classify as inexperienced.[26] Figure 6 compares the distribution of resale profits for experienced and inexperienced bidders.

[26] Over half of the volume in our matched data is accounted for by bidders who win just a single TM auction (overall, not just restricted to matched data). The remaining 27% of volume corresponds to bidders with 2-9 transactions. We did not find discernible differences between outcomes for bidders who win 1 auction versus those who win 2-9, so we do not break out their performance separately.

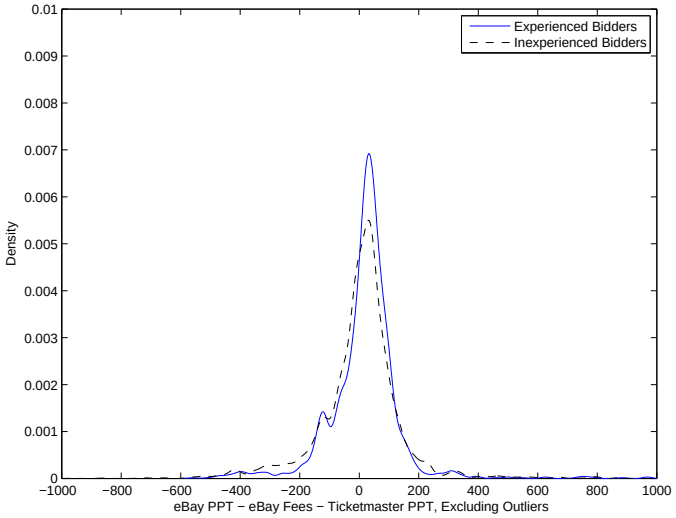


Figure 6: Resale profits of TM buyers with different levels of primary-market experience

As can be seen in Figure 6, experienced bidders perform better on average than do less experienced bidders. The difference in means is statistically significant, with experienced bidders purchasing tickets with resale profits of \$18.09 per ticket, while inexperienced bidders purchase tickets with resale profits of \$0.54 per ticket (p-value < 0.01).[27]

The figure also suggests that experience accounts for some of the asymmetry in the distribution of arbitrage profits. Specifically: (i) the most experienced bidders are significantly more likely to generate small positive profits of between \$0 and \$100 per ticket: 54.8% of transactions versus 42.6% of transactions (p-value < 0.01); and (ii) the most experienced bidders are significantly less likely to have large losses that exceed -\$100 per ticket: 11.6% of transactions versus 15.7% of transactions (p-value < 0.01). That is, the mode of small positive profits is disproportionately experienced bidders, whereas the fat left tail of large losses is disproportionately inexperienced bidders.

The fact that experienced bidders earn small positive profits on average is reassuring, because economic logic dictates that professional resellers should earn a return for the time and effort associated with reselling. While we cannot say whether \$18.09 per ticket is large or small relative to time and effort costs, we emphasize that it is an order of magnitude smaller than the \$134.94

[27] We performed a decomposition of the \$17.55 difference in profits between experienced and inexperienced bidders, and found that the difference is driven mostly by concert selection (\$5.42), and section-row selection (\$8.30). There are also differences from artist selection (\$1.02) and the ability to obtain seats within a section-row at a lower price than inexperienced bidders (\$2.80).

per ticket that bidders earn from resale under the counterfactual of using face values prices instead of the auction.

5.3 Resale Profits: Variation by Quality Level

Figure 7 depicts the distribution of arbitrage profits by ticket quality level. Since different auctions have different numbers of quality groups, we focus on tickets that are in the best- and worst-quality group in a particular auction, and also on tickets that are above and below the median quality group.

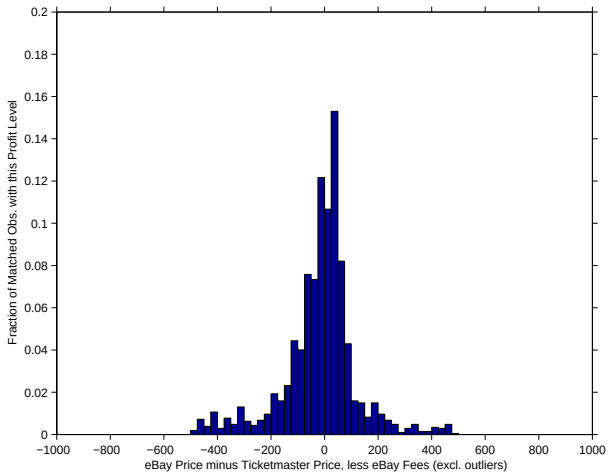
Interestingly, it is the highest quality group that has the lowest arbitrage profits, with a mean of -$15.62. This is driven by a fat left tail, and relates to a market-design proposal we will discuss in Section 6.2. The lowest quality group had strictly positive profits, of $21.23. However, when we look at above-median versus below-median quality groups the pattern reverses, with above-median groups enjoying positive profits ($18.17 on average) and below-median groups having losses on average (-$7.56).

5.4 Early vs. Late Resale

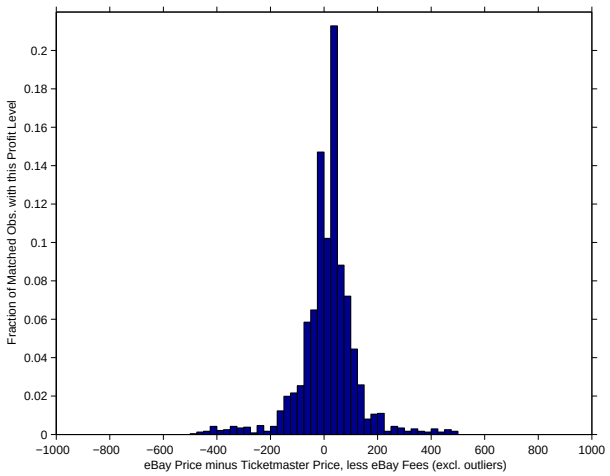
Our analysis above presumes that there is a single static aftermarket value for each concert-section-row, and uses eBay transaction data to estimate this aftermarket value. A recent paper by Sweeting [2012] suggests, however, that aftermarket prices evolve dynamically in a predictable manner. In particular, Sweeting finds that the secondary-market value of baseball tickets declines dramatically in the 30 days leading up to a game. This finding is consistent with dynamic pricing theory (cf. Board and Skrzypacz [2010]). Intuitively, sellers set a high price early, gradually lower the price, and ultimately attempt to salvage whatever revenue they can before the event passes and their tickets become worthless.

In this section we explore how the distribution of arbitrage profits varies when we divide the data based on the timing of the eBay aftermarket transaction. We classify a secondary-market transaction as "early" if it occurs in the first 30 days after the conclusion of the primary-market auction. We classify a transaction as "late" if it occurs in the last 30 days before the event occurs.[28]

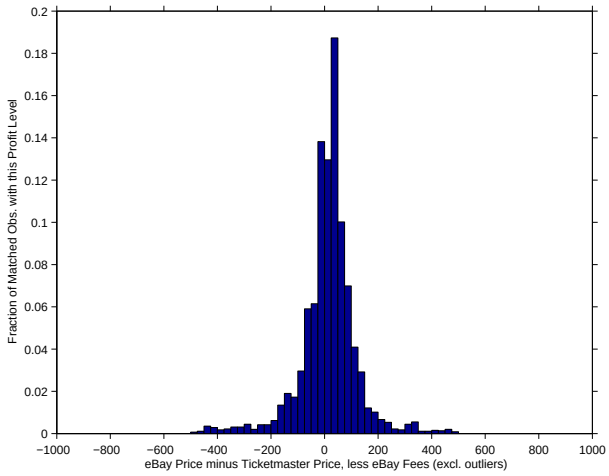
[28] If the number of days between the TM primary-market auction and the concert is $30 < t < 60$, then we define the early period as the first $t - 30$ days and the late period as the last $t - 30$ days. If the number of days between the TM primary-market auction and the concert is less than 30, we omit that concert from the test.



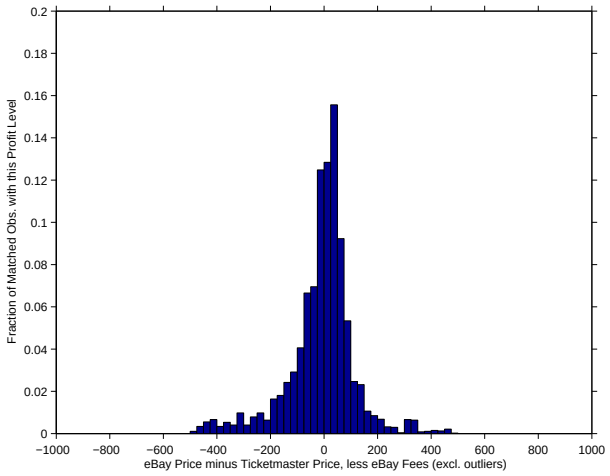
(a) Highest quality tickets



(b) Lowest quality tickets



(c) Above-median quality tickets



(d) Below-median quality tickets

Figure 7: Resale profits of tickets of varying quality

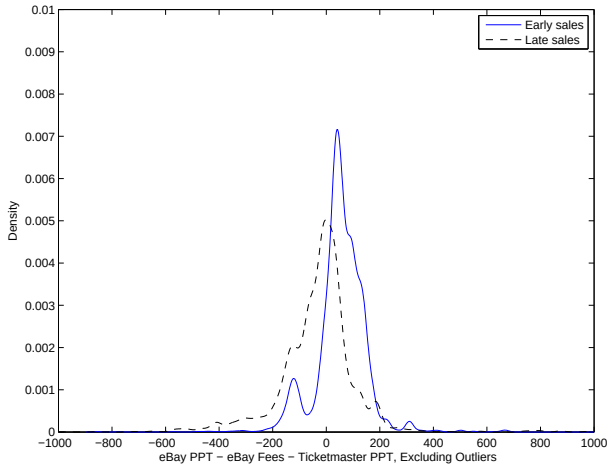


Figure 8: Resale profits for early and late secondary-market listings

Figure 8 compares the distribution of resale profits between early and late secondary-market listings.

Consistent with Sweeting [2012], there is a large and statistically significant difference between the two distributions, with the Early distribution having statistically significant positive profits, with a mean of $56.47 per ticket (95% CI: [$37.60, $73.25]), and the Late distribution having statistically significant negative profits, with a mean of -$33.06 per ticket (95% CI: [-$52.38, -$14.21]).[29] Analogously to the experienced versus inexperienced comparison, this is driven both by the Early distribution having a higher mode than the Late distribution ($42.56 for Early, -$0.99 for Late), and the Late distribution having a fatter left tail. Another interesting feature of the comparison of Early to Late eBay resales is that the primary-market auction prices are substantially more informative of Early resale values: the R^2 of TM primary-market auction prices on Early eBay prices is 0.806, versus 0.524 for Late, and 0.647 for the full sample.

5.5 Selection Concerns

Our eBay data are not a complete census of secondary-market activity, and for this reason one might worry about selection. We have three specific concerns.

The first potential concern is that our analysis only uses *successful* eBay listings – listings where either the seller's Buy-it-Now price was accepted, or where the seller's auction elicited bids of at least their reserve price. For any particular listing, our estimate of aftermarket value conditional on

[29] As we discuss in the next subsection, the large magnitude of the difference between Early and Late profits may be partly driven by selection.

success would be higher than that conditional on failure, so we might worry about a positive bias entering our analysis of the secondary market. This would cause us to *over*-estimate the prevalence of profitable resale opportunities. Since our main results suggest that arbitrage profits are either small or negative, we do not need to worry about this type of selection driving our main results.

A second potential concern is specific to eBay Buy-it-Now listings. Suppose that eBay sellers set their Buy-it-Now price equal to their tickets' true average aftermarket value plus a noise term that represents seller error. We then will observe more sales when the seller error term is negative than when it is positive. This will cause a negative bias in our estimate of secondary-market values, and hence cause us to *under*-estimate the returns to speculation. However, arbitrage profits are actually higher for Buy-it-Now listings (+$23.19) than other listings (-$16.66).[30] Thus, our results do not appear to be driven by eBay sellers who set their Buy-it-Now prices too low.

A third potential concern relates to us only seeing the portion of the aftermarket that occurs on eBay, as opposed to other venues such as StubHub. If the eBay component of the aftermarket is a random sample of the total aftermarket, then our use of just eBay data will cause us to have less precision than if we had the full aftermarket, but will not cause bias. If the eBay component is non-random, however, this could cause bias. The specific bias we are most worried about arises from the following. Suppose that a seller's strategy is to initially post their tickets at a high fixed price on a venue such as StubHub, and only if that fixed-price posting is unsuccessful, to run an auction on eBay. That is, what if sellers use eBay as a last-minute "salvage market" to ensure that their tickets are sold.[31] In this case eBay prices will be lower than average, which would cause us to under-estimate the returns to speculation. The large difference in arbitrage profits between Early and Late resales that we observed in Section 5.4, while broadly consistent with dynamic pricing theory and Sweeting [2012], suggests that this is an important concern.

One piece of evidence against this concern comes from looking at the difference between Early resales conducted using eBay's pure fixed-price selling format and eBay's pure auction format. Early BIN listings are associated with large, statistically significant profits of $66.87 per ticket.

[30] This finding is consistent with results in Einav et al. [2011] which suggest that Buy-it-Now prices are consistently higher than non-BIN prices, across a wide variety of eBay categories.

[31] The institutional reason why one should worry that sellers initially post on StubHub and then salvage on eBay, as opposed to the other way around, is the difference in the two venues' fee structures. StubHub does not charge listing fees and allows sellers to maintain their fixed price listing for as long as they like; eBay does charge listing fees, and depending on the type of seller most listings last for 7-10 days. Thus, posting a ticket at a high fixed price for a long period of time is free on StubHub (ex opportunity costs), but costly on eBay.

Early auctions, by contrast, are associated with negative profits of -$28.55 per ticket.[32] Early BIN listings are where we should be most worried about our first selection concern, namely, that we only observe arbitrage profits for successful eBay listings. Using a high BIN price early also is consistent with optimal dynamic pricing behavior, as discussed in Section 5.4. Early auction listings, on the other hand, should represent an unbiased estimate of the aftermarket value of the tickets at that particular moment in time. The fact that these profits are negative suggests that the TM primary-market auctions are not leaving large positive secondary-market profits on the table. A second piece of evidence against this concern is that Sweeting [2012] finds, in baseball ticket data from the same time period as our data, that posted prices are quite similar between eBay and StubHub.

Still, we find this selection concern worrying. A conservative response is to discard secondary-market data from the final days before the concert occurs, which is when salvage market activity is most likely. If we discard secondary market data from the final 30 days before the concert occurs, our estimate of mean arbitrage profits is $40.81 (95% CI: [$26.21, $54.95]). If we discard the final 15 days, the estimate is $34.67 (95% CI: [$20.59, $49.42]). We think of these estimates as a conservative upper bound on arbitrage profits, since they ignore that failing to resell early is a real risk in this market.

As one last test for selection concerns, Table 3 simply compares the set of concert-section-rows in our matched sample to the set of concert-section-rows in the full TM primary-market data set. While there are some differences between the two samples, the table suggests that our matched dataset is broadly representative of the tickets available in the TM primary-market auction.

6 Improving TM's Auction Design

In this section we discuss two potential improvements to TM's auction design, guided by the data.

[32] Following Einav et al. [2011], we classify an eBay sale as pure fixed price if the listing uses a Buy-it-Now price and does not allow bidders to bid less than the BIN amount, and classify an eBay sale as a pure auction if it does not use a BIN and uses a low starting bid. The -$28.55 figure in the text defines a low starting bid as <50% of the ticket's face value. If we use <10% instead the figure is $-26.81.

	Full TM Dataset	TM Data Matched with eBay	Difference
Average Price per Ticket ($)	308.76	277.49	-31.27
95% CI ($)	[278.99, 340.52]	[252.30, 302.97]	[-45.24, -18.12]
Average Markup per Ticket (%)	120.29	96.19	-24.11
95% CI (%)	[109.53, 132.09]	[85.61, 108.11]	[-31.65, -16.58]
Share of Top 5 artists (%)			
Rock & Roll Artist 4	19.90	21.05	1.15
Country Artist 1	8.37	10.28	1.92
Country Artist 3	8.03	8.67	0.63
Rock & Roll Artist 6	7.17	10.23	3.06
Rock & Roll Artist 5	6.59	6.43	-0.16
Share of Top 5 cities (%)			
New York	11.05	10.21	-0.84
Los Angeles	9.74	9.85	0.11
Boston	4.56	5.12	0.56
Chicago	4.23	3.75	-0.48
Miami	3.25	3.20	-0.06

Table 3: Comparison of summary statistics of full TM dataset and TM observations that match with eBay.

6.1 Increasing Bid Expressiveness

In mid-2007, TM increased the expressiveness of its auction's bidding language. Originally, bids consisted just of a dollar amount per ticket, and a desired number of tickets. After the rules change, bids could also include a minimum acceptable quality threshold. At the end of the auction, as before, bids are sorted in descending order, with the highest bid winning the highest-quality tickets, etc. The only modification is that if a bid is reached where the quality that would be awarded is below the bidder's threshold, then that bid is skipped.

The idea behind this rules change is to protect bidders from winning lower quality tickets than expected. Consider a bidder who bids $600 in the auction depicted in Figure 2. If the auction closed without any further bidding activity, he would win tickets in the first ticket group. However, if there are bids subsequent to his own, he might win tickets in any of the other ten ticket groups, or none at all. In principle, if a bidder wins tickets considerably worse than expected, the bidder could end up with negative realized surplus.[33] By allowing the bidder to attach a minimum acceptable

[33]In our stylized model (cf. Appendix B), this phenomenon can occur in equilibrium.

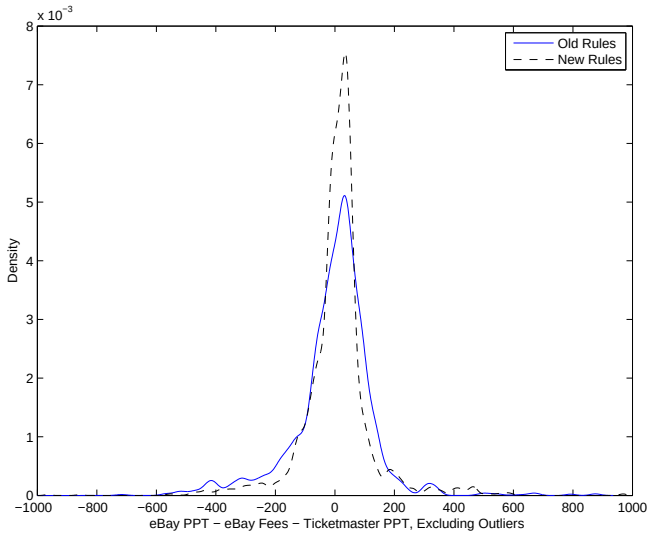


Figure 9: Secondary Market Profits: Old versus New Rules

quality to their bid, TM gives bidders a way to protect themselves against negative surplus.

Figure 9 displays the distribution of mean arbitrage profits under the old rules and the new rules. As can be seen from the figure, mean arbitrage profits are roughly similar under the two sets of auctions, but the *variance* of arbitrage profits changes. Specifically, the variance of arbitrage profits is lower under the new rules than the old rules. This result has the interpretation that the rules change leads to more accurate pricing. Intuitively, the quality threshold makes the TM auction a bit less like an all-pay auction – in which bidders pay a certain amount for uncertain quality, and hence the distribution of equilibrium bids is compressed (cf. Bulow and Levin, 2006) – and a bit more like a standard multi-object auction, in which bidders can bid different amounts for different quality goods.

A natural follow-on to this first increase in expressiveness is to allow bidders to express different valuations for different quality levels. For instance, in the auction depicted in Figure 2, TM might allow a bidder to express that he is willing to pay $600 for tickets in the first 3 rows, or $400 for tickets in rows 4-10.

6.2 Reducing Strategic Complexity

TM's auction design is simple to explain but complex to play. Unlike many modern auction designs, which are either exactly strategyproof or satisfy a large-market notion of strategyproofness (cf.

Azevedo and Budish [2012], optimal play of the TM auction design depends heavily on one's beliefs about the behavior of other auction participants.[34] For this reason, we should worry about the possibility of bidder mistakes. This is especially so given the consumer nature of the goods being auctioned.

From the consumer's perspective, the most harmful mistake is submitting a bid that is substantially higher than is necessary to win even the highest-quality tickets. The intuition is as follows. Suppose a bidder has a type in the middle of the value distribution, such that in equilibrium he should win a middle-quality ticket, say the 5th row. If he bids a bit too high, he will pay more but get a better quality ticket (say, 3rd row). If he bids a bit too low, he will pay less and get a lower quality ticket (say, 7th row), or pay nothing and get nothing. The resale market protects the bidder from both of these kinds of mistakes: even though he receives the incorrect ticket given his preferences, he pays roughly the correct price for it, and can reoptimize using the resale market if he so wishes. In the terminology of Roth [2008], the game is "safe" for this kind of bidder, even though the game is not strategyproof. Suppose instead that the bidder has a very high value, such that in equilibrium he should get one of the highest-quality seats, in say the 1st row. If he bids too little, he will pay less and get lower quality seats (say, 3rd row), and can reoptimize in the resale market. If, however, he bids too high, he will still get assigned tickets in the 1st row, but will now pay substantially above the true market value for these tickets. The auction is "unsafe" for such a bidder.

Consistent with this worry, we find evidence of occasional severe overbidding for tickets in the highest quality row in a given auction (e.g., 1st row): 5.9% of winning bids for tickets in the best row pay at least 50% more than would have been necessary for seats in that row, and 3.2% of winning bids are at least 100% higher than necessary. These overbids are rarely submitted by experienced bidders, and are disproportionately submitted by first-time bidders, who then disproportionately never return to bid in another auction. See Table 4. Thus, while overbidding might increase an artist's revenue in a particular auction, it has a detectable negative effect on the long-run health of the TM auction marketplace. Additionally, strategic complexity might have a negative effect on the health of the marketplace which we cannot detect in our data, namely that some potential

[34] In the equilibrium we derive in Appendix B, bidders bid the probability-weighted value of the ticket they receive, conditional on the event that they win any ticket at all, less a shading term that represents a probability-weighted information rent. While this model is stylized, and in particular abstracts from the information bidders learn over the course of the auction, it highlights the complexity of the strategic environment.

Cutoff Overbid Percentage	50%	100%	Share in Full Dataset
% overbids at least this large	5.91%	3.18%	
Of the overbids			
% experienced bidders	1.39%	2.07%	14.81%
% inexperienced bidders	98.61%	97.93%	85.19%
% for whom it was only win	71.03%	65.80%	58.43%
% for whom it was first win	84.96%	82.38%	68.26%
% for whom it was last win	80.78%	76.68%	68.94%
% who bid on last day	58.77%	76.68%	67.41%
% who bid in last hour	45.13%	63.21%	53.94%
% who bid in last minute	19.22%	28.50%	19.96%

Table 4: Overbidding analysis. See the text for description.

bidders never participate in the first place.

Fortunately, this is an auction design problem which, once diagnosed, can be fixed. In particular, this is the same issue Milton Friedman observed in US Treasury auctions, and which motivated Friedman's proposal of Uniform Price Auctions to replace Pay-as-Bid Auctions: in a Uniform Price Auction "no one is deterred from bidding by fear of being stuck with an excessively high price" [Friedman, 1991]. We describe two specific auction design alternatives. The first, and more straightforward, possibility would be for TM to modify its payment rule in a manner analogous to moving from Generalized First Price to Generalized Second Price. Specificially, all winners of tickets in the 1st row pay the highest bid amongst winners of the 2nd row; all winners of tickets in the 2nd row pay the highest bid amongst winners of the 3rd row; etc. This is exactly like GSP except that there are multiple units available at each quality level. Hence, an appropriate name might be the Generalized Uniform Price Auction.

The Generalized Uniform Price Auction is not strategyproof, even in a large-market sense, for the same reason that GSP is not strategyproof (cf. Edelman et al. [2007]). A second possibility, then, would be to move to a fully strategyproof auction design. If we assume that each bidder is single-minded about the quantity of tickets they require – that is, a bidder who requires 4 tickets derives no value from obtaining 2 tickets, and hence is not tempted to shade his demand – then a proxy version of the simultaneous ascending auction is strategyproof in this context.[35]

[35]See Kelso and Crawford [1982] and Demange et al. [1986] for early studies of simultaneous ascending auctions, and see Ausubel and Milgrom [2002] for a proxy version that would be strategyproof in this context. Without the single-mindedness assumption, the simultaneous ascending auction is not strategyproof, but it would satisfy the large-market notion of approximate strategyproofness defined in Azevedo and Budish [2012]. As the market grows

A disadvantage of the simultaneous ascending auction, relative to GUPA, is that it requires more preference information. Specifically, the auction would need to elicit each bidder's value for each row, or block of rows. Given this information, the SAA works by having each bidder's proxy bid on the row that, at current prices, maximizes their surplus. For instance, if a bidder's value is \$1000 for the 1st row and \$500 for the 2nd row, and current prices are \$800 for the 1st row and \$400 for the second row, the proxy would bid on the first row, since $\$1000 - \$800 = \$200$ is more surplus than $\$500 - \$400 = \$100$. At the end of the auction, this bidding process yields the minimum competitive equilibrium price vector with respect to the reported demands (to within a linear function of the bidding increment, see Milgrom [2000]), and hence is efficient by the first welfare theorem.

7 Conclusion

This paper studies Ticketmaster's introduction of auctions into the primary market for event tickets. Our basic findings suggest that the auctions work (as auctions should!): price discovery is improved; artist revenues more than double versus the fixed-price counterfactual; and, perhaps most importantly, the auctions eliminate or at least substantially reduce resale profits for speculators. This last finding suggests that auctions can eliminate much of the socially wasteful rent-seeking activity that has been associated with this market since at least the 19th century. In the simplest model of rent-seeking competition, in which rents get fully competed away by rent-seeking activities, the \$9mm of incremental revenue generated by Ticketmaster's auctions in our data set represents a commensurate improvement in social welfare.

Yet, at least as of the present writing, auctions remain a small component of the primary market for event tickets. Why has adoption not been more rapid? One possible explanation highlighted within our analysis is the strategic complexity of the TM auction design. We conclude by discussing two explanations that are outside the scope of our analysis.

First, the market has developed a viable substitute for auctions, specifically, setting posted prices in a more sophisticated way. While it would be difficult to estimate the market-clearing price for tickets to one-of-a-kind events such as Charles Dickens' reading tour, there are many entertainment events for which it is possible to develop a reasonable estimate of market-clearing

large, the probability that shading one's quantity demanded from 4 tickets to 2 affects market-clearing prices in an advantageous way vanishes to zero.

prices. For instance, major league baseball teams host 80+ home games per year, and can use historical secondary-market data and sales patterns to estimate demand for a particular game.[36] If demand information is sufficiently precise and the number of tickets being sold is large, optimally-set posted prices should closely approximate the efficiency, revenue and no-arbitrage performance of an auction.

Second, artists frequently express an aversion to setting market-clearing prices for their tickets, often citing concerns for fairness and accessibility.[37] A particularly interesting episode in this regard is the 2007 Miley Cyrus / Hannah Montana tour mentioned in the introduction. This tour used primary-market auctions for some of its earliest concert dates, but not for later dates, perhaps due to the negative publicity associated with high ticket prices.[38] For the artist's next tour, in 2009, Disney again set below-market prices, but this time adopted technology that essentially eliminates the possibility of secondary-market activity: just as airplane tickets are non-transferable because they are attached to the passenger's name, 2009 Miley Cyrus / Hannah Montana tickets were non-transferable because they were attached to a specific credit card, that had to be presented in person at the concert venue [Waddell, 2009]. In our view, turning off the secondary market is a sensible substitute for setting market-clearing prices in the primary market: in either case, rent-seeking activity by speculators is eliminated.

References

Lawrence Ausubel and Paul Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1), 2002.

Eduardo Azevedo and Eric Budish. Strategyproofness in the large. *Working Paper*, 2012.

Patrick Bajari and Ali Hortaçsu. The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34(2):329–355, 2003.

[36] The industry jargon for this practice is "dynamic pricing". For recent press coverage see Goldstein [2012] and Rishe [2012].

[37] See Che et al. [2011] and Condorelli [2012] for formal models that can be used to understand this phenomenon.

[38] A second and similar example is Next Restaurant in Chicago, a popular restaurant which uses tickets instead of reservations. Asked why the restaurant does not use auctions to set market-clearing prices (the tickets sell out very quickly and there is an active resale market), the proprietor wrote: "Simply, we never want to invert the value proposition so that customers are paying a premium that is disproportionate to the amount of food / quality of service they receive." [Baliga, 2011].

Patrick Bajari and Ali Hortaçsu. Economic insights from internet auctions. *Journal of Economic Literature*, 42(2):457–486., 2004.

Sandeep Baliga. How to increase revenue at Next restaurant, April 2011. URL `http://cheaptalk.org/2011/04/11/how-to-increase-revenue-at-next-restaurant/`.

Gary S. Becker. A note on restaurant pricing and other examples of social influences on price. *Journal of Political Economy*, 99(5):1109–1116, October 1991.

Simon Board and Andy Skrzypacz. Revenue management with forward looking buyers. *Working Paper*, 2010.

Jeremy Bulow and Paul Klemperer. The generalized war of attrition. *American Economic Review*, 89(1):175–189, March 1999.

Jeremy Bulow and Jonathan Levin. Matching and price competition. *American Economic Review*, 96(3):652–668, June 2006.

Yeon-Koo Che, Ian Gale, and Jinwoo Kim. Assigning resources to budget-constrained agents. Working Paper, July 2011.

Daniele Condorelli. Market and non-market mechanisms for the optimal allocation of scarce resources. Working Paper, February 2012.

Marie Connolly and Alan Krueger. Rockonomics: The economics of popular music. *Handbook on the Economics of Art and Culture*, Chapter 20:667–719, 2006.

Pascal Courty. An economic guide to ticket pricing in the entertainment industry. *Louvain Economic Review*, 2000.

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

Pascal Courty. Ticket pricing under demand uncertainty. *Journal of Law and Economics*, XLVI(2): 627–52, 2003b.

Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–72, 1986.

Economist. Bring on the touts. *Economist*, August 4, 2012.

Benjamin Edelman and Michael Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43(1):192–198, February 2007.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.

Liran Einav, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. Learning from seller experiments in online markets. *NBER Working Paper No. 17802*, 2011.

Milton Friedman. How to sell government securities. *Wall Street Journal*, August:A8, 1991.

Jacob Goldstein. Baseball teams are acting like airlines, March 2012. URL `http://www.npr.org/blogs/money/2012/03/05/147982995/at-the-ballpark-ticket-prices-change-as-fast-as-the-weather?`

Nathan Hubbard. An open letter to fans and partners on ticket reselling, August 2011. URL `http://blog.ticketmaster.com/2011/08/`.

Alexander Kelso and Vincent Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.

Steven E. Landsburg. *The Armchair Economist: Economics and Everyday Life*. The Free Press, 1993.

Phillip Leslie and Alan Sorensen. The welfare effects of ticket resale. NBER Working Paper No. 15476, November 2009.

Steven D. Levitt. What do I have in common with Hannah Montana?, January 2008. URL `http://www.freakonomics.com/2008/01/08/what-do-i-have-in-common-with-hannah-montana/`.

Mitch Lipka. Paperless tickets: Is Ticketmaster hurting consumers?, March 2011. URL `http://blogs.reuters.com/reuters-money/2011/03/29/paperless-tickets-is-ticketmaster-hurting-consumers/`.

Ticketmaster Entertainment LLC. Form 10-K, February 2010.

Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272, 2000.

Paul Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, September 1982.

Julie Holland Mortimer, Chris Nosko, and Alan Sorensen. Supply responses to digital distribution: Recorded music and live performances. *Information Economics and Policy*, 24(1):3–14, March 2012.

Sucharita Mulpuru. The future of online secondary ticketing. *Forrester Research Report*, 2008.

Chris Nelson. Media; ticketmaster auction will let highest bidder set concert prices, September 2003. URL `http://www.nytimes.com/2003/09/01/business/media-ticketmaster-auction-will-let-highest-bidder-set-concert-prices.html?pagewanted=all`.

New York (State) Department of Law. Why can't I get tickets?: Report on ticket distribution practices / Office of New York State Attorney General Eliot Spitzer, 1999.

New York Times Editorial Page. A carnival of crime, November 1876.

Patrick Rishe. Dynamic pricing: The future of ticket pricing in sports, January 2012. URL `http://www.forbes.com/sites/prishe/2012/01/06/dynamic-pricing-the-future-of-ticket-pricing-in-sports/`.

Ellen Rosen. In the race to buy concert tickets, fans keep losing. *New York Times*, October 6, 2007.

Sherwin Rosen and Andrew M. Rosenfield. Ticket pricing. *Journal of Law and Economics*, 40: 351–376, October 1997.

Alvin E. Roth. Repugnance as a constraint on markets. *Journal of Economic Perspectives*, 21(3): 37–58, Summer 2007.

Alvin E. Roth. What have we learned from market design? *Economic Journal*, 118(March): 285–310, 2008.

Alvin E. Roth and Axel Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *American Economic Review*, 92(4):1093–1103, September 2002.

Kerry Segrave. *Ticket Scalping: An American History, 1850-2005*. McFarland & Co., 2007.

Ethan Smith and Sarah Silver. To protect its box-office turf, ticketmaster plays rivals' tune. *Wall Street Journal*, September 12, 2006.

Steve Stecklow. Concerts: Can't get no...tickets? *Wall St*, January 7, 2006.

Andrew Sweeting. Price dynamics in perishable goods markets: The case of secondary markets for major league baseball tickets. *Journal of Political Economy*, Conditionally accepted., 2012.

New York Times. Mr. dickens' readings – sale of tickets for the second course, Dec 1867a.

New York Times. Mr. dickens in boston - the eages demand for tickets, Nov 1867b.

North America TM Team. The ongoing battle to make access to tickets fair for you, June 2011. URL `http://blog.ticketmaster.com/2011/06/`.

Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.

Ray Waddell. Miley Cyrus fights scalpers with paperless tickets, June 2009. URL `http://www.reuters.com/article/2009/06/21/us-cyrus-idUSTRE55K0CX20090621`.

Kim Zetter. Wiseguys plead guilty in Ticketmaster captcha case, November 2010. URL `http://www.wired.com/threatlevel/2010/11/wiseguys-plead-guilty/`.

Unit of Analysis	N	Gross/Net Profits	Avg. Profit ($)	Clustered CI ($)	Mode	Skewness
TM Transaction	$7,973$	Net	4.27	$[-10.31, 16.71]$	33.48	-0.59
		Gross	15.28	$[0.47, 27.85]$	40.97	-0.39
CSR Combination	$1,591$	Net	-17.26	$[-26.29, -8.68]$	27.32	-1.51
		Gross	-7.25	$[-16.33, 1.42]$	32.88	-1.30
eBay Transaction	$3,035$	Net	-15.12	$[-27.55, -3.42]$	29.36	-0.70
		Gross	-5.27	$[-18.04, 6.66]$	34.24	-0.50

Table 5: Arbitrage Profits from Buying in TM Primary-Market Auctions and Reselling on eBay. Net profits are calculated by subtracting eBay fees. Confidence intervals are calculated at the 95% level. Standard errors are clustered at the concert level and are calculated via bootstrap.

A Robustness Tests

A.1 Robustness to Alternative Matching Specifications

Table 5 lists the average profits associated with buying tickets in the TM primary-market auctions and then reselling in the eBay secondary market, for each of the matching specifications discussed in Section 4.3. We report gross profits as well as profits net of eBay transaction fees, which are roughly 4% of the eBay sale price. 95% confidence intervals are calculated using the bootstrap. We construct these confidence intervals with observations clustered at the concert level, because it seems natural to think of each concert as its own market. Although some tickets for some concerts are sold in multiple auctions, we expect unobservables such as secondary market demand to be correlated within a concert.

In our main specification, our estimate of average arbitrage profits is $15.28 before fees, and $4.27 net of fees. Over all of our specifications, the 95% confidence intervals admit estimates of gross arbitrage profits ranging from -16.33 to +27.85, and of net arbitrage profits ranging from -26.29 to +16.71. These figures suggest that the average profits associated with buying in the TM primary-market auction and then reselling in the eBay secondary market are economically small.

A.2 Sub-Sample Robustness

In this section we explore the robustness of the no-arbitrage result on various sub-samples of our data.

A.2.1 Resale Profits: Variation by Artist

Figure 10 depicts the distribution of resale profits for the 6 artists in our data set for whom we are able to match at least 500 TM transactions.

Some artists have systematically negative resale profits (e.g., Country Artist 1), while others have systematically positive resale profits (e.g., Hip Hop / R&B Artist 4). Our basic takeaway from Figure 10, however, is the overall robustness of the no arbitrage result as first depicted in Figure 1.

A.2.2 Resale Profits: Variation by City

Figure 11 depicts the distribution of resale profits for the 5 cities in our data that were visited by more than 10 of the artists in our data.

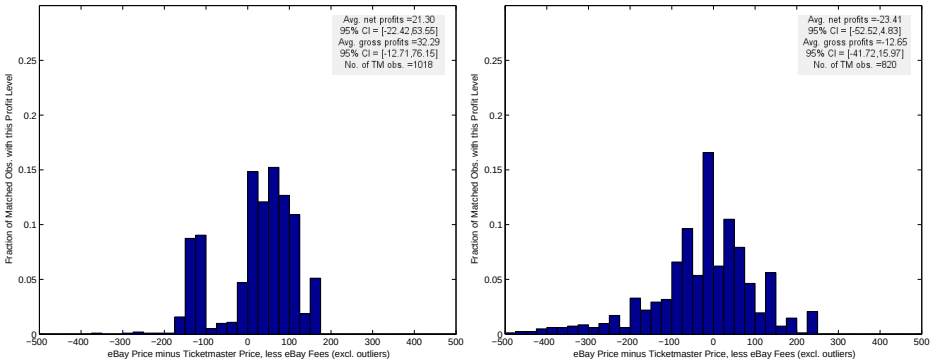
Our main takeaway from the decomposition by city as depicted by Figure 11 is similar to our takeaway from the decomposition by artists, above: the no arbitrage pattern first depicted in 1 is robust when examined on subsets of the full data set. It is interesting as well that the basic shape of Figure 1, where the distribution of arbitrage profits has a positive mode, but a fatter left tail than right tail, is so prevalent across cities.

B Model

In this Appendix, we provide a theoretical analysis of the Ticketmaster primary-market auction design, focusing on the original 2003 auction rules. Our analysis both clarifies the relationship between the TM auction design and the position auctions used widely in internet advertising markets, and formalizes that the TM auction design is "sensible" in that it satisfies attractive efficiency, revenue and no arbitrage properties.

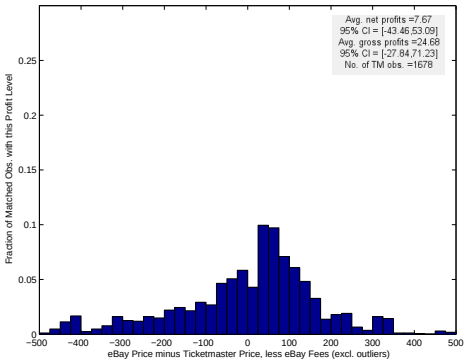
Our model closely mirrors that of Edelman et al. [2007]. There are K (pairs of) tickets and $N > K$ ex-ante symmetric bidders.[39] Initially, we think of a bidder as either a fan who intends to use the tickets himself, or as a speculator acting as a proxy agent on behalf of a specific fan. Below we will endogenize entry by speculators, in a stylized manner, to derive a simple no-arbitrage result. The tickets are vertically differentiated, with bidder i's private valuation for the k^{th} best ticket equal

[39] For expositional purposes, we will refer to a pair of tickets simply as one object, or one ticket.

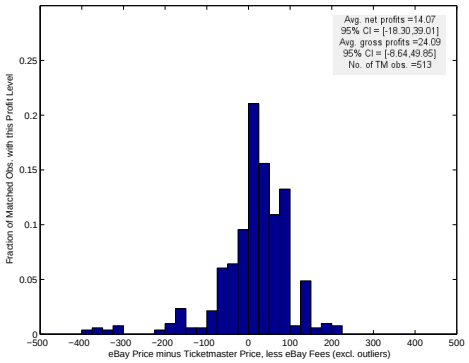


(a) Hip Hop / R&B Artist 4

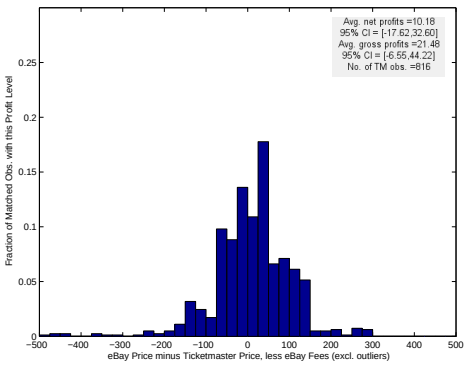
(b) Country Artist 1



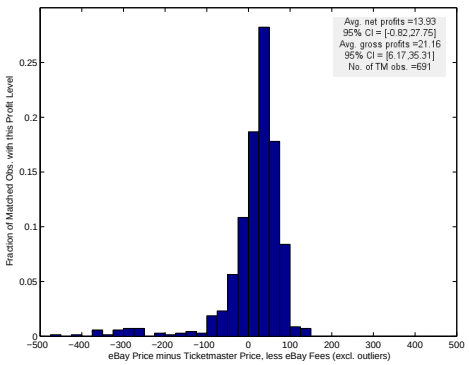
(c) Rock & Roll Artist 4



(d) Rock & Roll Artist 5

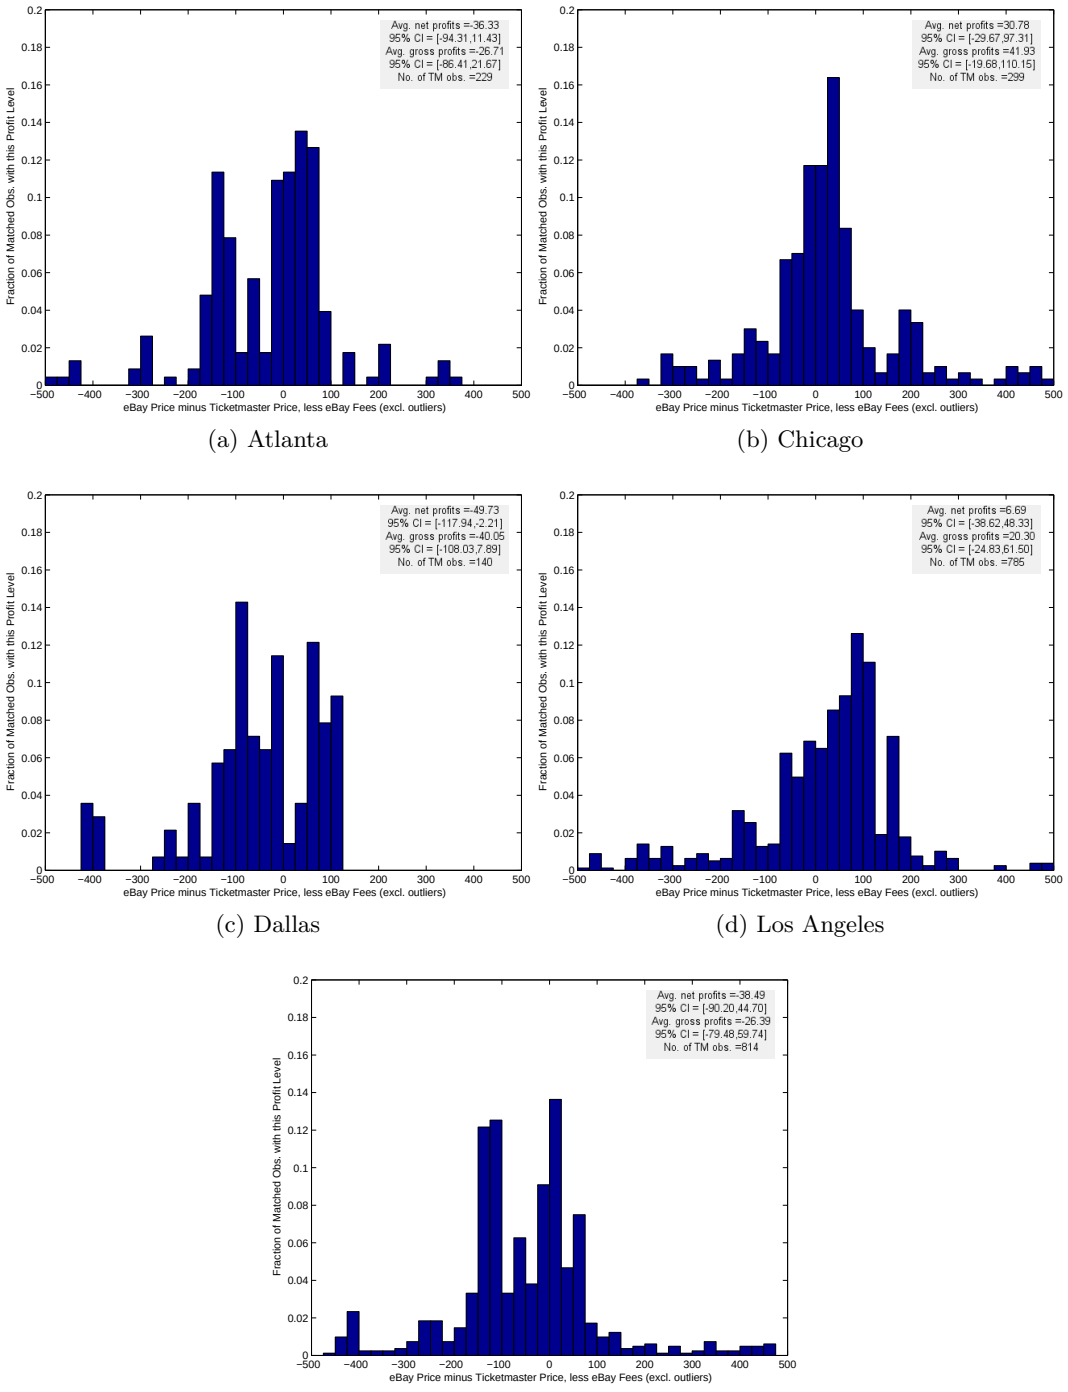


(e) Rock & Roll Artist 6



(f) Country Artist 3

Figure 10: Resale profits by artist



(a) Atlanta

(b) Chicago

(c) Dallas

(d) Los Angeles

(e) New York

Figure 11: Resale profits by city, for cities visited by at least 10 artists. In this figure, cities are aggregated at the level of Metropolitan Statistical Area.

to $\alpha_k v_i$: v_i is bidder i's type, and α_k describes the quality of ticket k, with $\alpha_1 > \alpha_2 > \cdots > \alpha_K$.[40] We assume that the quality-levels are common knowledge, and we normalize $\alpha_K = 1$.

Each bidder's type v_i is drawn independently and identically from a distribution with cdf $F(\cdot)$ and support $[0, \bar{v}]$. We assume that $F(\cdot)$ is continuously differentiable, with $f(\cdot)$ the corresponding pdf. The distribution of preferences is common knowledge.

We consider two ways to model the Ticketmaster auction design, sealed-bid and ascending auction, analogously to how Edelman et al. [2007] consider both a sealed-bid and ascending auction variant of GSP. The sealed-bid model captures the fact that the TM auction uses a "hard-close" ending rule [cf. Roth and Ockenfels, 2002], and highlights that bidders are uncertain, at the moment they bid, of what quality ticket they will win, and indeed whether they will win any ticket at all. However, the TM auction is not static, and in particular bidders do have some information at the time they bid about the demand of other bidders (see Figure 2 above). Hence we also consider an ascending auction model, for completeness. Our main theory results – on efficiency, constrained revenue maximization, and no arbitrage – obtain under both models.

In the *sealed-bid TM auction* model, each bidder submits a single bid. The bids are ranked in descending order, and then the k^{th} highest bidder wins the k^{th} ticket, for each $k = 1, \ldots, K$. There is no reserve price. Winning bidders pay their bid amount, while losers pay nothing. To explain the difference between this model and Edelman et al. [2007]'s model of GSP, let $b_{(k)}$ denote the k^{th}-highest bid, for some $k \leq K$. In GSP, this bidder's total payment is $b_{(k+1)}\alpha_k$: the next highest bid, times the click-through rate. In our model, this bidder's total payment is simply $b_{(k)}$: his own bid, without any adjustment for the realized quality.

Our *ascending TM auction* model is related to the Generalized English Auction (GEA) of Edelman et al. [2007], analogously to how our sealed-bid model is related to their treatment of GSP. An auction clock, initialized at $p = 0$, ascends continuously at the rate of \$1 per unit time. Bidders can "drop" out of the auction at any time; once a bidder drops out of the auction, the auction is over for him [cf. Milgrom and Weber, 1982]. The auction ends when all bidders but one have dropped. The last remaining bidder gets the best ticket, and pays the amount at which the next-to-last bidder dropped. The k^{th}-to-last remaining bidder, for $k = 2, \ldots, K$, gets the k^{th} ticket, and pays the amount at which he himself dropped. The $N - K$ bidders who do not get a

[40] In Edelman et al. [2007], α_k is the click-through rate of the kth slot, and v_i is the ith advertiser's private value per click.

ticket pay zero.

B.1 Efficiency and Revenue Results

B.1.1 Efficiency

Equilibrium of the Sealed-bid TM Auction Let $P_k(v)$ denote the probability that a bidder whose type is v has the k^{th} highest type out of the n bidders. We show the following.

Proposition 1. *There exists a symmetric monotonic Bayes-Nash equilibrium of the sealed-bid TM auction in which all bidders bid according to*

$$b(v) = \frac{1}{\sum_{k=1}^{K} P_k(v)} \left(\sum_{k=1}^{K} P_k(v) \left(v\alpha_k\right) - \sum_{k=1}^{K} \int_0^v \alpha_k P_k(x) dx \right). \tag{1}$$

The resulting allocation is efficient.

Function (1) can be interpreted as follows. The first term, $\frac{\sum_{k=1}^{K} P_k(v)(v\alpha_k)}{\sum_{k=1}^{K} P_k(v)}$ is bidder v's expectation of the value of the ticket she will receive, conditional on being one of the k winners. Note that this term will be strictly between the value she places on the best ticket, $\alpha_1 v$, and the value she places on the worst ticket, $\alpha_K v$. The second term, $\frac{\sum_{k=1}^{K} \int_0^v \alpha_k P_k(x) dx}{\sum_{k=1}^{K} P_k(v)}$, is the amount by which she shades her bid due to the pay-as-bid nature of the auction. If $K = 1$ this is just the standard single-unit auction information rent. When $K > 1$, the numerator places relatively more weight, in determining how much to shade, on tickets that are of high quality (the α_k term) and on tickets where bidder v's value is high enough that it is likely that someone with a lower value than she wins those tickets (the $\int_0^v P_k(x)dx$ term). Intuitively, if a ticket is of very low quality (low α_k), or if bidder v is not really in the running for the ticket (the $\int_0^v P_k(x)dx$ term), then she should not earn an information rent from that ticket.

Equilibrium of the Ascending TM Auction Consider a bidder of type v. Let $\underline{v}$ be the lowest possible type who has not dropped out when all other bidders are following symmetric equilibrium strategies. At a given point in time, aside from the bidder in question, suppose that there are k other active bidders in the auction. Then let $T^O\left(v; \underline{v}, k\right)$ denote the amount of time bidder v would be willing to wait before dropping out, *as long as none of the other active bidders drop out in this duration*. Additionally, define the hazard rate in the standard manner, $h\left(v\right) = \frac{f(v)}{1-F(v)}$.

Proposition 2. *The unique symmetric perfect Bayesian equilibrium of the ascending auction is defined by,*

$$T^O(v;\underline{v},k) = \begin{cases} v-\underline{v}, & \text{if } k \geq K, \\ (\alpha_k - \alpha_{k+1}) \int_{\underline{v}}^{v} xkh(x)\,dx, & \text{if } k < K. \end{cases} \tag{2}$$

The resulting allocation is efficient.

The equilibrium of the ascending auction can be understood as follows. Given a bidder with value v, as long as there are at least K other active bidders in the auction she behaves as if she is competing in a standard $K+1^{st}$-price auction for K units of ticket K. That is, she simply bids up to her value for the K^{th} ticket of $v\alpha_K \equiv v$, i.e., her waiting time is $v-\underline{v}$. Once ticket K has been allocated the game changes in an important way: now, bidder v behaves as if she is competing in an *all-pay* auction against $K-1$ other bidders for the quality increment $\alpha_K - \alpha_{K-1}$. That is, she is competing for the right *not* to wind up with ticket K. The all-pay nature of the auction follows from the fact that waiting is now costly: since she is a winner in the auction, she must pay her bid. If she survives this auction, she competes against the $K-2$ other remaining bidders in an all-pay auction for the quality increment $\alpha_{K-1} - \alpha_{K-2}$, and so forth. The intuition for the equilibrium waiting time is as in Bulow and Klemperer [1999]: bidders equate their marginal cost of waiting, $\frac{\partial T^O(v;\underline{v},k)}{\partial v}$, with their marginal benefit from doing so, $k(\alpha_k - \alpha_{k+1})\,vkh(v)$. In the Appendix, we prove by induction that the above collection of individual auction equilibria constitutes an equilibrium of the full ascending TM auction game.

B.1.2 Revenue

By Myerson's Lemma [cf. Milgrom, 2004], since the equilibria described in Section B.1.1 lead to an efficient allocation, we immediately have the following corollary:

Corollary 1. *The sealed-bid and ascending Ticketmaster auctions are revenue equivalent to any other efficient auction design, such as Generalized Second Price or Vickrey-Clarke-Groves.*

B.2 No Arbitrage

So far our analysis has treated the set of bidders as exogenous; these bidders were conceptualized either as fans or as professional resellers acting as proxies on behalf of specific fans.

In this section we allow for endogenous entry by professional resellers, in the following stylized manner. There is a continuum of potential bidders in the population, of which a fraction β are professional resellers, and the remaining are fans. Fix $\varepsilon > 0$ and $w \in (\varepsilon, \bar{v})$. Fans' types are drawn independently and identically from the continuously differentiable distribution $F_{fan}(\cdot)$ with support $[0, \bar{v}]$. Each reseller's type is drawn independently and identically from the continuously differentiable distribution $F_{pro}(\cdot)$ with support $[w - \varepsilon, w]$. The interpretation is that w is the expected price (per unit quality) of a ticket on the secondary market. Resellers also face a small idiosyncratic cost, in the interval $[0, \varepsilon]$, associated with participating in the aftermarket. The purpose of the idiosyncratic cost is to ensure that $F_{pro}(\cdot)$ can be assumed to be atomless.

In each auction, n bidders are randomly drawn from the above population of potential bidders. Notice that this process is equivalent to taking n i.i.d. draws from the distribution $F(\cdot)$, where

$$F(x) = \beta F_{pro}(x) + (1 - \beta) F_{fan}(x).$$

$F(\cdot)$ has support on $[0, \bar{v}]$, and we assume that $F(\cdot)$ is continuously differentiable $\forall \beta \in [0, 1]$.

We then model entry by speculators by increasing β, the share of speculators in the population from which the bidders are drawn.[41] This approach to endogenizing entry allows for the following simple no arbitrage statement:

Proposition 3. *Let $s(v; \beta)$ denote the expected surplus of bidder v when the bidders are drawn from a population that has a fraction β of professional resellers. Under either the sealed-bid or ascending TM auction, $\forall v \in [w - \varepsilon, w], \lim_{\beta \to 1^-} s(v; \beta) < \varepsilon \alpha_1$.*

Proposition 3 formalizes in a simple manner that, under free entry by speculators, we should expect speculators to earn negligible resale profits. Though a simple result, it highlights an important difference between auctions and fixed-price selling mechanisms, which is that auctions eliminate resale profits so long as there is sufficient entry.

[41] The number of bidders n can also be increased, but this is not necessary for the result. Proposition 3 is expressed for constant n, but the result holds (and in fact speculator surplus is smaller) if n is allowed to increase with β.

B.3 Proofs

Proof of Proposition 1

Proof. Let us look for a symmetric equilibrium in which all bidders use the same bidding function $b(\cdot)$. We initially assume and later prove that $b(\cdot)$ is strictly increasing, so that there is a one-to-one relationship between bids and valuations. Therefore we can think of a bid $\hat{b}$ as the submission of a valuation $\hat{v}$, such that $\hat{b} = b(\hat{v})$. The bidder can then be thought of as choosing the submitted valuation $\hat{v}$ optimally, given the bidding function $b(\cdot)$. The bidder thus maximizes her expected value from the auction by solving the following program:

$$\max_{\hat{v}} \sum_{k=1}^{K} \left[v\alpha_k - b(\hat{v})\right] P_k(\hat{v}),$$

where

$$P_k(x) = \binom{n-1}{k-1} F(x)^{n-k} \left(1 - F(x)\right)^{k-1} \tag{3}$$

is the probability that a bidder with valuation x wins the k^{th} object. In order for $b(\cdot)$ to define a symmetric equilibrium, the first order condition requires that the bidder's expected value must be maximized at her true valuation v. That is,

$$\sum_{k=1}^{K} \left(\left[v\alpha_k - b(v)\right] P_k'(v) - P_k(v)b'(v)\right) = 0. \tag{4}$$

We can use 4 to solve for the equilibrium bidding function as follows. Rearranging terms,

$$\begin{aligned} \sum_{k=1}^{K} v\alpha_k P_k'(v) &= \sum_{k=1}^{K} \left(b(v)P_k'(v) + P_k(v)b'(v)\right) \\ &= \frac{d}{dv}\left(b(v)\sum_{k=1}^{K} P_k(v)\right) \\ \implies b(v)\sum_{k=1}^{K} P_k(v) - b(0)\sum_{k=1}^{K} P_k(0) &= \int_0^v x \sum_{k=1}^{K} \alpha_k P_k'(x)dx. \end{aligned}$$

But $P_k(0) = 0, \forall k \in \{1, \cdots, K\}$, and so

$$b(v) = \frac{1}{\sum_{k=1}^{K} P_k(v)} \int_0^v x \sum_{k=1}^{K} \alpha_k P_k'(x) dx. \tag{5}$$

We can then use integration by parts to derive 1:

$$\begin{aligned} b(v) &= \frac{1}{\sum_{k=1}^{K} P_k(v)} \sum_{k=1}^{K} \int_0^v x \alpha_k P_k'(x) dx \\ &= \frac{1}{\sum_{k=1}^{K} P_k(v)} \left(\sum_{k=1}^{K} \alpha_k \left[x P_k(x) \right]_{x=0}^{x=v} - \sum_{k=1}^{K} \int_0^v \alpha_k P_k(x) dx \right) \\ &= \frac{1}{\sum_{k=1}^{K} P_k(v)} \left(\sum_{k=1}^{K} P_k(v) \left(v \alpha_k \right) - \sum_{k=1}^{K} \int_0^v \alpha_k P_k(x) dx \right). \end{aligned}$$

Finally we need to show that the bidding function is strictly increasing, since the previous analysis was predicated on this premise. Notice that

$$b'(v) = \frac{\left(\sum_{k=1}^{K} P_k(v) \right) \left(v \sum_{k=1}^{K} \alpha_k P_k'(v) \right) - \left(\sum_{k=1}^{K} P_k'(v) \right) \left(\int_0^v x \sum_{k=1}^{K} \alpha_k P_k'(x) dx \right)}{\left(\sum_{k=1}^{K} P_k(v) \right)^2}$$

$$\implies Sign\left\{ b'(v) \right\} = Sign\left\{ \left(\sum_{k=1}^{K} P_k(v) \right) \left(v \sum_{k=1}^{K} \alpha_k P_k'(v) \right) - \left(\sum_{k=1}^{K} P_k'(v) \right) \left(\int_0^v x \sum_{k=1}^{K} \alpha_k P_k'(x) dx \right) \right\}.$$

Lemma 1. For any $x \in [0, \bar{v}]$, $\sum_{k=1}^{K} \alpha_k P_k'(x) > 0$.

Proof. Let $x \in [0, \bar{v}]$. Then[42]

[42] For expositional purposes, we define $\alpha_{K+1} \equiv 0$.

$$
\begin{aligned}
\sum_{k=1}^{K} \alpha_k P_k'(x) &= \sum_{k=1}^{K} \alpha_k \binom{n-1}{k-1} F(x)^{n-k-1}(1-F(x))^{k-2}\left[(n-k)(1-F(x))-(k-1)F(x)\right] f(x) \\
&= (n-1)f(x)\left(\sum_{k=1}^{K} \alpha_k \binom{n-2}{k-1} F(x)^{n-k-1}(1-F(x))^{k-1} - \sum_{k=2}^{K} \alpha_k \binom{n-2}{k-2} F(x)^{n-k}(1-F(x))^{k-2}\right) \\
&= (n-1)f(x)\left(\sum_{k=1}^{K} \alpha_k \binom{n-2}{k-1} F(x)^{n-k-1}(1-F(x))^{k-1} - \sum_{k=1}^{K} \alpha_{k+1} \binom{n-2}{k-1} F(x)^{n-k-1}(1-F(x))^{k-1}\right) \\
&= \sum_{k=1}^{K} \left[\alpha_k - \alpha_{k+1}\right](n-1)\binom{n-2}{k-1} F(x)^{n-k-1}(1-F(x))^{k-1} f(x) \\
&> 0.
\end{aligned}
$$

The first equality follow from the definition of $P_k'(\cdot)$. The second equality is obtained by shifting the index of the second sum and using $\alpha_{K+1} = 0$. The last equality results from grouping terms, and the inequality is due to the fact that $\{\alpha_k\}_{k=1}^{K+1}$ is a strictly decreasing sequence. This completes the proof of the Lemma. □

From Lemma 1,

$$
\begin{aligned}
\int_0^v x \sum_{k=1}^{K} \alpha_k P_k'(x)dx &< v \int_0^v \sum_{k=1}^{K} \alpha_k P_k'(x)dx \\
&= v \sum_{k=1}^{K} \alpha_k P_k(v).
\end{aligned}
$$

Hence

$$
\begin{aligned}
&\left(\sum_{k=1}^{K} P_k(v)\right)\left(v \sum_{k=1}^{K} \alpha_k P_k'(v)\right) - \left(\sum_{k=1}^{K} P_k'(v)\right)\left(\int_0^v x \sum_{k=1}^{K} \alpha_k P_k'(x)dx\right) \\
&> v\left[\left(\sum_{k=1}^{K} P_k(v)\right)\left(\sum_{k=1}^{K} \alpha_k P_k'(v)\right) - \left(\sum_{k=1}^{K} P_k'(v)\right)\left(\sum_{k=1}^{K} \alpha_k P_k(v)\right)\right].
\end{aligned}
$$

Therefore it suffices to show that

$$\left(\sum_{k=1}^{K} P_k(v)\right)\left(\sum_{k=1}^{K} \alpha_k P_k'(v)\right) - \left(\sum_{k=1}^{K} P_k'(v)\right)\left(\sum_{k=1}^{K} \alpha_k P_k(v)\right) \geq 0$$

$$\iff \sum_{j=1}^{K}\sum_{k=1}^{K} (\alpha_j - \alpha_k)\, P_k(v) P_j'(v) \geq 0$$

$$\iff \sum_{j=1}^{K-1}\sum_{k=j+1}^{K} \underbrace{(\alpha_j - \alpha_k)}_{>0} \Big[P_k(v)P_j'(v) - P_j(v)P_k'(v)\Big] \geq 0.$$

The second and third inequalities follows from rearranging the terms on the left hand sides of the first and second inequalities, respectively. It would then be sufficient to show that $\forall v \in [0, \bar{v}], \forall j, k$ such that $j < k \leq K, P_k(v)P_j'(v) - P_j(v)P_k'(v) \geq 0$. Note that

$$\begin{aligned}
&P_k(v)P_j'(v) \\
&= \left[\binom{n-1}{k-1} F(v)^{n-k}(1-F(v))^{k-1}\right]\left[\binom{n-1}{j-1} F(v)^{n-j-1}(1-F(v))^{j-2}\left[(n-j)(1-F(v)) - (j-1)F(v)\right] f(v)\right] \\
&= \binom{n-1}{k-1}\binom{n-1}{j-1} F(v)^{2n-j-k-1}(1-F(v))^{j+k-3}\left[(n-j) - (n-1)F(v)\right] f(v).
\end{aligned}$$

The first equality is due to the definitions of $P_k(\cdot)$ and $P_j'(\cdot)$ and the second is obtained from algebraic simplification. It follows that

$$\begin{aligned}
P_k(v)P_j'(v) - P_j(v)P_k'(v) &= \binom{n-1}{k-1}\binom{n-1}{j-1} F(v)^{2n-j-k-1}(1-F(v))^{j+k-3}\left[(n-j) - (n-k)\right] f(v) \\
&= (k-j)\binom{n-1}{k-1}\binom{n-1}{j-1} F(v)^{2n-j-k-1}(1-F(v))^{j+k-3} f(v) \\
&\geq 0.
\end{aligned}$$

Therefore the bidding function $b(\cdot)$ is indeed strictly increasing, which implies that the resulting allocation is efficient. This completes the proof. □

Proof of Proposition 2

Proof. The proof will proceed by induction. We will construct K auctions that, if held sequentially, yield the GEA. First, a second-price English auction is conducted among the N bidders for K units

of object K, each worth v_i to bidder i (recall that $\alpha_K = 1$). The last K surviving bidders win the objects, and pay the price at which the last drop-out occurred. Bidders draw their private valuations from $F(\cdot)$, and the mechanics of the auction are as described in Section B.1.1.

Once the above auction has concluded, the K winners of the auction immediately enter an all-pay English auction in which K bidders compete for $K-1$ upgrades from object K to object $K-1$, each worth $(\alpha_{K-1} - \alpha_K) v_i$ to bidder i. After one drop-out, the remaining bidders win the object, and all bidders (including the losing bidder) pay the price at which the drop-out occurred. Again, the mechanics of the auction as are described in Section B.1.1. Notice that at the end of this auction, the $K-1$ winners have object $K-1$, the loser keeps object K, and the total payment across the two auctions of all K participants is the drop-out price of the losing bidder in the auction for the upgrade to object $K-1$.

Another all-pay auction is then held among the $K-1$ winners of the previous auction, offering $K-2$ upgrades from object $K-1$ to object $K-2$. Proceeding in this manner, a total of $K-1$ all-pay auctions are conducted sequentially among the winners of the previous auction. In each such auction, an upgrade of one quality level is awarded to all bidders but one. The process naturally stops after the quality upgrade from item 2 to item 1 has been sold.

Losing bidders pay nothing, and all winning bidders pay the price at which they drop out (except for the winner of object 1, who pays the price at which the second-to-last survivor drops out). Therefore when the above K auctions are conducted sequentially, the resulting composite auction is the GEA.

Now let us analyze any of the above-described all-pay auctions in isolation: $k+1$ bidders compete for k objects, each worth $(\alpha_k - \alpha_{k+1}) v_i$ to bidder i. The following Lemma will be useful in our proof:

Lemma 2. *The unique symmetric perfect Bayesian equilibrium of the above all-pay English auction is characterized by:*

$$T^O(v; \underline{v}, k) = k(\alpha_k - \alpha_{k+1}) \int_{\underline{v}}^{v} xh(x)\, dx.$$

Proof. This result follows immediately from Lemma 3 of Bulow and Klemperer [1999]. □

Consider also the initial second-price auction for object K in isolation.

Lemma 3. *The unique symmetric perfect Bayesian equilibrium of the above second-price English auction is characterized by:*

$$T^O(v; \underline{v}, K) = v - \underline{v}.$$

Proof. Since losing bidders do not pay, a bidder v would strictly prefer to remain in the auction if the total time that has elapsed is less than v, since winning object K would give her positive surplus in this event. Conversely, she would strictly prefer to drop out if the total time that has elapsed is greater than v, because she would get negative surplus if she were to win the object. Therefore she will drop out when the total time that has elapsed is v. When all bidders follow this strategy, her waiting time at any instant can be written as $T^O(v; \underline{v}, k) = v - \underline{v}$. □

In what follows, we will refer to the subgame that begins with the auction for the upgrade to object k as subgame k. Now analyze subgame 1, from the perspective of a bidder with valuation v. Recall that this is the last subgame of the GEA. There is one other bidder remaining, and the bidders are competing for one upgrade from object 2 to object 1. Both bidders have already won object 2, but importantly the benefit from this object is sunk. Likewise their waiting costs from surviving to this point in the game are sunk. Therefore we can conclude from Lemma 2 that the unique symmetric perfect Bayesian equilibrium of this subgame is defined by (2).

Next, consider subgame k, $1 < k < K$. Consider a bidder with value v, and suppose that all other players bid according to (2). Suppose additionally that in the each subgame $j < k$, the bidder in question knows that she will follow the proposed strategies, conditional on surviving until subgame j is reached. We know from Lemma 2 that (2) gives the myopic best response of the bidder in the auction for the upgrade to object k. Hence it could not be optimal to deviate by dropping out earlier than she would under the proposed equilibrium - the bidder would be giving up positive expected utility in the present auction and possibly in future auctions.

Suppose instead that the bidder drops out later than she would under (2), i.e., she plays as if she has valuation $v^* > v$. There are three possibilities, each of which occur with positive probability:

1. Both types v^* and v would not win object k. Then the bidder strictly prefers to bid as type v, since she would have a lower drop-out price by doing so.

2. Both types v^* and v would win object k. Then since the bidder follows $T^O(v; \cdot, \cdot)$ in all future auctions, her expected payoff is the same from playing as either type.

3. Type v would drop out in the auction for object k, but type v^* would win object k. At the start of subgame $k-1$, v is lower than the lowest possible type $\underline{v}_{k-1}$who has not dropped out. Since the bidder follows $T^O(v; \underline{v}_{k-1}, k-1)$ in this auction, she must drop out immediately. Therefore her expected utility from misrepresenting her type differs from that under truthful play only in terms of her payoffs from the auction for object k. But from Lemma 2, her expected surplus in this auction is maximized by playing truthfully.

Thus the bidder strictly prefers to follow (2) instead of bidding as a type with a higher valuation. It follows that when all other bidders follow the proposed equilibrium strategies and the bidder knows that she will follow (2) in all future subgames, her optimal waiting time in the auction for object k is also given by $T^O(v; \cdot, k)$. By induction, then, (2) defines the unique symmetric perfect Bayesian equilibrium of subgame k, $k < K$.

Moving to the first auction, consider again a bidder with value v, who knows that all other bidders are bidding according to (2), and that she will also follow the proposed equilibrium strategies in future subgames $k < K$, if she is still active in those subgames. When all other players follow the proposed strategies, they drop out precisely when the price equals their valuation of object K. Lemma (3) tells us that this strategy is the myopic best response in the auction for object K. Using an argument analogous to the one above, the bidder cannot profitably deviate from (2) by playing as though her valuation is higher or lower than v. By an additional step of induction, we can thus conclude that (2) defines the unique symmetric perfect Bayesian equilibrium of the GEA.

Finally, the fact that the equilibrium is efficient follows immediately from (2); $\forall k, T^O(\cdot; \underline{v}, k)$ is clearly strictly increasing. This completes the proof. □

Proof of Proposition 3

Proof. First notice that since the sealed-bid and ascending TM auctions are efficient, any bidder's expected surplus in either auction can be expressed as the expected surplus in the sealed-bid

auction, which equals:

$$\begin{aligned} s(v;\beta) &= \sum_{k=1}^{K} \left[v\alpha_k - b(v)\right] P_k(v;\beta) \\ &= \sum_{k=1}^{K} \int_0^v \alpha_k P_k(x;\beta)\, dx, \end{aligned}$$

where the second equality uses (1), and both expressions make explicit the dependence of the probability of winning object k on the share of professional resellers. It suffices to show that $\forall \delta > 0$, $\exists \mu > 0$ such that $1 - \beta < \mu \implies s(v;\beta) < m + \delta$, $\forall v \in [w - \varepsilon, w]$, where $m < \varepsilon\alpha_1$.

Let $\delta > 0$ and $v \in [w - \varepsilon, w]$.We can write

$$\begin{aligned} s(v;\beta) &= \sum_{k=1}^{K} \int_0^v \alpha_k P_k(x;\beta) dx \\ &= \int_0^{w-\varepsilon} \sum_{k=1}^{K} \alpha_k P_k(x;\beta) dx + \int_{w-\varepsilon}^{v} \sum_{k=1}^{K} \alpha_k P_k(x;\beta) dx \\ &< \alpha_1 \left[\int_0^{w-\varepsilon} \sum_{k=1}^{K} P_k(x;\beta) dx + \int_{w-\varepsilon}^{v} \sum_{k=1}^{K} P_k(x;\beta) dx \right]. \end{aligned}$$

The second equality follows from splitting the integral, and the inequality uses the fact that $\{\alpha_k\}_{k=1}^{K}$ is a strictly decreasing sequence.

Define $m \equiv \alpha_1 \int_{w-\varepsilon}^{v} \sum_{k=1}^{K} P_k(x;\beta) dx$. Since $\sum_{k=1}^{K} P_k(x;\beta) \leq 1$, with strict inequality when $x < w$, $m < \alpha_1 \int_{w-\varepsilon}^{v} 1 \cdot dx = (v - w + \varepsilon)\alpha_1 \leq \varepsilon\alpha_1$. Therefore

$$\begin{aligned} s(v;\beta) &< m + \alpha_1 \int_0^{w-\varepsilon} \sum_{k=1}^{K} P_k(x;\beta)\, dx \\ &< m + \alpha_1 (w - \varepsilon) \sum_{k=1}^{K} P_k(w - \varepsilon;\beta), \end{aligned}$$

The second inequality uses the fact that the probability of winning some object, $\sum_{k=1}^{K} P_k(\cdot;\beta)$, is strictly increasing.

Now, choose $\mu = \min\left\{\frac{n}{K} - 1, \left(\frac{\delta}{\alpha_1 K\binom{n-1}{K-1}(w-\varepsilon)}\right)^{\frac{1}{n-K}}\right\} > 0$. Then

$$\begin{aligned}\alpha_1(w-\varepsilon)\sum_{k=1}^{K} P_k(w-\varepsilon;\beta) &= \alpha_1(w-\varepsilon)\sum_{k=1}^{K}\binom{n-1}{k-1}F(w-\varepsilon;\beta)^{n-k}(1-F(w-\varepsilon;\beta))^{k-1} \\ &\leq \alpha_1(w-\varepsilon)\sum_{k=1}^{K}\underbrace{\binom{n-1}{k-1}F(w-\varepsilon;\beta)^{n-k}}_{\equiv G(w-\varepsilon;k,\beta)}.\end{aligned}$$

The equality uses (3) and the inequality is due to the fact that $1 - F(w-\varepsilon;\beta) \leq 1$. Consider any $k \in \{1, 2, \cdots, K-1\}$. Then

$$\begin{aligned}\frac{G(w-\varepsilon;k,\beta)}{G(w-\varepsilon;k+1,\beta)} &= \frac{\binom{n-1}{k-1}F(w-\varepsilon;\beta)^{n-k}}{\binom{n-1}{k}F(w-\varepsilon;\beta)^{n-k-1}} \\ &= \frac{kF(w-\varepsilon;\beta)}{n-k} \\ &= \frac{k}{n-k}\left[\beta F_{pro}(w-\varepsilon) + (1-\beta)F_{fan}(w-\varepsilon)\right] \\ &= \frac{k}{n-k}(1-\beta)F_{fan}(w-\varepsilon) \\ &< \frac{k}{n-k}\mu F_{fan}(w-\varepsilon) \\ &< \frac{k}{n-k}\mu \\ &\leq 1, \forall k \leq K.\end{aligned}$$

The second equality follow from algebraic manipulations, the third uses the definition of $F(\cdot;\cdot)$ and the fourth uses the fact that $F_{pro}(w-\varepsilon) = 0$. The first inequality is due to the fact that $1 - \beta < \mu$ (by assumption) and the second uses $F_{fan}(w-\varepsilon) < 1$. The last (weak) inequality results because $\forall k \leq K, \frac{k}{n-k} \leq \frac{K}{n-K}$, and $\mu \leq \frac{n}{K} - 1 \implies \frac{K}{n-K}\mu \leq 1$. Therefore $\{G(w-\varepsilon;k,\beta)\}_{k=1}^{K}$ is a strictly increasing sequence, and so

$$\begin{aligned}
\alpha_1(w-\varepsilon)\sum_{k=1}^{K} G(w-\varepsilon;k,\beta) &< \alpha_1(w-\varepsilon)\,KG(w-\varepsilon;K,\beta) \\
&= \alpha_1(w-\varepsilon)\,K\binom{n-1}{K-1}F(w-\varepsilon;\beta)^{n-K} \\
&= \alpha_1(w-\varepsilon)\,K\binom{n-1}{K-1}\left[\beta F_{pro}(w-\varepsilon)+(1-\beta)F_{fan}(w-\varepsilon)\right]^{n-K} \\
&= \alpha_1(w-\varepsilon)\,K\binom{n-1}{K-1}\left[(1-\beta)F_{fan}(w-\varepsilon)\right]^{n-K} \\
&< \alpha_1(w-\varepsilon)\,K\binom{n-1}{K-1}\mu^{n-K}.
\end{aligned}$$

The second equality follows from the definition of $G(\cdot;\cdot,\cdot)$, the third uses the definition of $F(\cdot;\cdot)$ and the fourth uses the fact that $F_{pro}(w-\varepsilon)=0$. The inequality uses $F_{fan}(w-\varepsilon)<1$ and $1-\beta<\mu$. But from the definition of μ,

$$\begin{aligned}
\mu &\leq \left(\frac{\delta}{\alpha_1(w-\varepsilon)\,K\binom{n-1}{K-1}}\right)^{\frac{1}{n-K}} \\
\implies \alpha_1(w-\varepsilon)\,K\binom{n-1}{K-1}\mu^{n-K} &< \delta \\
\implies s(v;\beta) &< m+\delta,
\end{aligned}$$

as required. The second inequality follows from algebraic manipulation, and the last follows from the chain of inequalities that has been developed throughout this proof. This completes the proof. □