

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

Aditya Bhave[†] and Eric Budish[‡]

January 26, 2014

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. What about using an auction to set price correctly? This paper studies the recent introduction of auctions into the event-ticket 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”: the auctions substantially improve price discovery, roughly double performer revenues, and, on average, nearly eliminate the arbitrage profits associated with underpriced tickets. The data thus suggest that auctions can eliminate the speculator rent-seeking that has been associated with this market since the 19th century, and that seems to have exploded in volume in the 21st century.

KEYWORDS: market design, auctions, primary markets, resale markets, rent-seeking, position auctions, event tickets, e-commerce

*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, Sandeep Baliga, Hoyt Bleakley, Peter Coles, Ben Edelman, Jeff Ely, Matt Gentzkow, Philip Haile, Ali Hortacsu, Paul Klemperer, Alan Krueger, Julie Mortimer, Roger Myerson, Chris Nosko, Matt Notowidigdo, Ariel Pakes, David Parkes, Canice Prendergast, Phil Reny, Al Roth, Michael Schwarz, Jesse Shapiro, Alan Sorensen, Lars Stole, Chad Syverson, Steve Tadelis, and Hal Varian for helpful discussions, as well as seminar participants at eBay Research, the University of Chicago, Duke, the 2012 NYU Stern IO Day, Stanford, the 2012 NBER Market Design Working Group Meeting, the 2012 Michigan Auction and Matching Workshop, and ACM EC 2013. We thank Yi Sun, 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 National Science Foundation (ICES-1216083) and 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 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 a performer, 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, 2013]. 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 \$2,000 [Levitt, 2008].

Economists have long found this phenomenon puzzling (Becker [1991], Landsburg [1993], Rosen and Rosenfield [1997], Courty [2003a], Baliga [2011]). First, underpricing reduces revenues for the performer. Second, underpricing encourages socially wasteful rent seeking by ticket speculators.¹

¹It is possible to tell stories for why performers may genuinely wish to sell tickets to their *fans* at a below-market price – e.g., social-good consumption complementarities [Becker, 1991], altruism (Che et al., 2013, though see also Bulow and Klemperer, 2012), or the sale of complementary goods over the artist’s lifecycle [Mortimer et al., 2012] – but it is hard to argue that artists genuinely wish for ticket speculators to make arbitrage profits. In that sense the

Modern information technology has only exacerbated the scale of the rent seeking. With both the primary and secondary markets almost exclusively online, what used to be a localized, labor-intensive activity in the pre-internet era now has few or no geographical boundaries and significant scale economies.² Recent industry estimates are that fully 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 In the colorful words of the Arkansas Attorney General, “All hell broke loose with Hannah Montana” [Rosen, 2007].

A natural question is: *what about using 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 scope for speculator rent seeking. 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]

true puzzle is the *combination* of low prices and an active secondary market. There have been some recent efforts by artists to set low prices but also eliminate the secondary market by making 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]. See further discussion on these issues in the conclusion.

²In the Dickens era, and as recently as the late 20th century, the basic rent-seeking technology in the primary market was getting a good spot in a physical queue, and much of the secondary market occurred outside the physical venue. A 1999 New York State Attorney General report described “diggers” in the primary market – groups that “push and intimidate their way to the front of the line” – and “scalpers” in the secondary market – individuals who stand “in front of or near the venue for which tickets are being sold” (New York Attorney General, 1999). In the present day, tickets can be amassed in the primary market online using software bots (and low-wage overseas workers who complete “captchas”) and sold in the secondary market on websites such as eBay and StubHub (Zetter, 2010; Ticketmaster Corporate Blog, June 2011).

³The 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 [2013] find that the rate of resale was on the order of 5% in summer 2004. Ticketmaster 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.”

This paper shows that TM’s primary-market auctions “worked”: the auctions substantially improved price discovery, roughly doubled artist revenues, and, on average, nearly eliminated the arbitrage profits between the primary market and the secondary market.

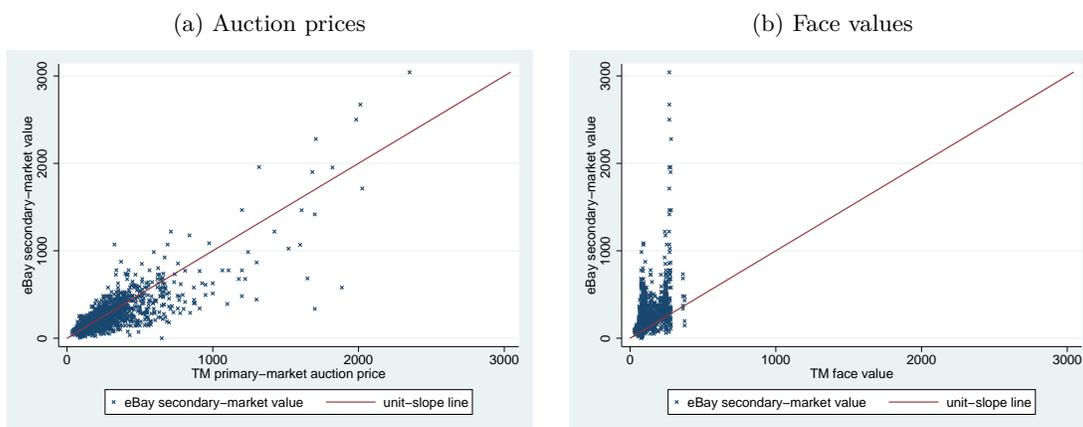
Our empirical research design is very simple. We have proprietary data from TM, which indicate the price at which each ticket was sold in TM’s primary-market auction, for all 2007 concert tours that used auctions (we will describe the auction rules in detail below, which are interesting in their own right). The data cover 22 distinct concert tours and 576 distinct concerts. We also have data from TM that 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 quality 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.

Figure 1 presents a scatterplot of our matched TM-eBay dataset, and conveys our main results. The left panel plots the TM primary-market auction price and the subsequent eBay secondary-market resale value. The data mostly cluster along the 45-degree line, which shows that the primary-market auction price is, on average, an accurate reflection of secondary-market value. The mean difference between the auction price and the subsequent resale value is just \$6.07, or around 2% of the average primary-market auction price of \$274. This difference is economically small and statistically indistinguishable from zero. In the right panel, instead of using the primary-market auction price, we use the ticket’s face value. Now, most of the data are above the 45-degree line, sometimes dramatically so – this is the underpricing phenomenon. The mean difference between the face value and the secondary-market resale value is \$136, or just under 100% of the average

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

⁵The 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 was standard practice in the secondary market at the time of our data. See Section 4.3 for more details.

Figure 1: Relationship between eBay secondary-market values and TM primary-market prices, both auction prices and face values



Notes: Panel (a) indicates the average primary-market auction price and secondary-market resale value for each concert-section-row in our matched dataset. Panel (b) indicates the ticket face value and average secondary-market resale value for each concert-section-row in our matched dataset. Prices are on a per-ticket basis. eBay secondary-market values are net of eBay fees. For more details on the data, see the text of Section 4.

primary-market face value of \$145.⁶ Moreover, the face-value prices contain much less information about secondary-market values than do the auction prices. In a regression of secondary-market value on primary-market prices, the R^2 is 0.66 using auction prices versus 0.24 using face-value prices. In sum, the auctions discover significantly higher prices than the counterfactual face values, and these prices are essentially correct on average.

We also explore the auction performance of the most experienced bidders in the TM auction – the top 1% of bidders, who account for 16% of volume. We find that the experienced bidders – the “Bob the Brokers” – do in fact statistically outperform the inexperienced bidders, in the sense that they purchase tickets in the auction with greater subsequent resale profits. However, the magnitude of their profits is still relatively modest, at \$19 per ticket, which is an order of magnitude smaller than the \$136 mean rent associated with purchasing primary-market tickets at their counterfactual face values. This \$19 per ticket can perhaps be interpreted as a return for the time, effort, and risk associated with ticket speculation (cf. Courty [2003a,b]).

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

⁶The large magnitude is consistent with Leslie and Sorensen [2013]’s finding that the most severe underpricing occurs for high-quality tickets, which are the focus of TM’s auctions.

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.). A potentially unappealing aspect of TM’s specific position auction design is its payment rule: in TM’s auction, bidders pay their own bid, whereas both theory and the observed evolution of the keyword advertising market suggest that it is more attractive to have a winning bidder’s payment depend only on bids lower than their own. In the terminology of Edelman et al. [2007], TM’s auction is more like a generalized first-price auction (GFP) than a generalized second-price auction (GSP). While we do not find evidence of the severe problems with GFP auctions documented in the keyword advertising setting [Edelman and Ostrovsky, 2007]⁷ – as described above, our results suggest that TM’s auctions work quite well overall – we do find evidence that bidders in TM’s auctions make occasional large bidding errors associated with the pay-as-bid nature of the auction, and that these errors are concentrated amongst inexperienced market participants. This suggests that there may be gains from modifying TM’s auction format to be less strategically complex.

We briefly mention three ways in which our results contribute to the broader market design literature beyond the specific context of event-ticket markets. First, our paper is, to our knowledge, the first empirical illustration of the usefulness of position auctions in a context other than online advertising, and also documents a novel variant on position auctions. These findings should be of interest to the literature on position auctions, which has been extremely active since Edelman et al. [2007] and Varian [2007]’s studies of position auctions for keyword advertising. Second, our results provide empirical support for the proposition that, in markets with resale, sensibly designed primary-market auctions accurately discover secondary-market resale values. This finding may be a useful input to market design debates in other contexts, for instance, the debate over whether to use auctions in the market for initial public offerings (IPOs). While there are of course many differences between concert tickets and shares of stock, there are some important similarities: (i) both are non-trivial to price; (ii) both have histories of severe underpricing; (iii) both have histories of elaborate

⁷One key difference between our setting and the keyword auction setting is that TM auctions are one-shot rather than repeated. Thus, the local instability of generalized first-price auction equilibria is not as problematic here as in the keyword auction setting. A second difference is that the difference in quality between successive prizes is frequently quite small. We thank Michael Schwarz for a helpful conversation on these points.

rent-seeking behavior associated with this underpricing (for IPOs, see, e.g., Nocera [2013]); and (iv) both have secondary markets that are widely viewed to be efficient, suggesting that accurate pricing in the primary market is a realistic possibility. Third, our results serve as a case study in the use of market design to reduce rent seeking. Perhaps the oldest objection to market design is to invoke the Coase theorem: market design details do not ultimately matter, because private trade will eventually lead to the socially optimal allocation.⁸ Our study is a reminder that this argument is flawed – even in the absence of Myerson-Satterthwaite bargaining frictions – because bad market design can induce socially wasteful rent-seeking behavior on the way to the ultimate allocation.

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 examines the role of bidder experience and discusses potential modifications to TM’s 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 fiscal year 2012, Ticketmaster sold more than 148 million event tickets valued at on the order of \$10 billion, on behalf of clients including venues, promoters, sports leagues and teams, and museums and cultural institutions.⁹ Tickets are typically sold at fixed prices that vary coarsely with seat quality, e.g., there might be just 3 or 4 pricing tiers in a venue with tens of thousands of seats. Tickets typically go on sale months in advance of an event.

“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

⁸Milgrom [2004], Section 1.4.1, calls this “One of the most frequent and misguided criticisms of modern auction design.”

⁹Ticketmaster merged with Live Nation, a promoter, venue operator and artist management firm, in January 2010. The 148 million tickets figure is from Live Nation’s 2012 annual report [Live Nation, 2013]. Live Nation also sold an additional 108 million tickets through its clients’ box offices. The \$10 billion estimate extrapolates from a volume figure of \$8.0 billion reported in Ticketmaster’s last financial filing as a standalone company, its 2009 annual report, based on 14% cumulative growth in the number of tickets sold (130 million tickets were reported sold in its 2009 report, [Ticketmaster, 2010]) and 7% growth in the consumer price index per the BLS.

the order of \$4bn annually, and other estimates range as high as \$10bn annually [Mulpuru, 2008, Stecklow, 2006]. At the time of our data, eBay was the largest forum for secondary-market activity [Mulpuru, 2008]; at present the largest forum is StubHub, which has grown substantially since eBay’s acquisition of it in 2007. 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]. See Sweeting [2012] for a detailed study of the secondary market, some findings from which manifest in our data as well; see further discussion in Section 5.2.

An important recent study by Leslie and Sorensen [2013] examines the welfare effects of the secondary market, and finds empirical evidence of substantial costs and benefits of 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 [2013]’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.¹⁰ There are at least three reasons. First, the internet has lowered the costs of amassing tickets in the primary market.¹¹ Second, the internet has lowered the cost of reselling tickets in the 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]).

¹⁰See 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.

¹¹For 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 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].

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.

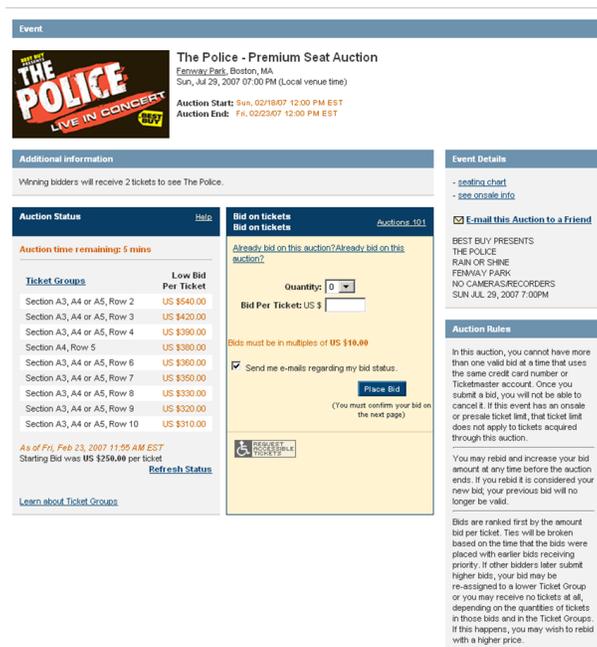
Which Tickets are Auctioned? For any particular event, the determination of which tickets to sell by auction (if any), and which to sell by fixed price, is made by TM’s client. In our data, an average of about 97 tickets are auctioned per concert, with a maximum of 862 tickets. The auctioned tickets 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”. This decision to focus on high-quality tickets is consistent with Leslie and Sorensen [2013]’s finding that high-quality tickets are associated with 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.

Auction Rules The auction itself lasts for several days, starting and ending at pre-announced fixed times. 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 opened on a Sunday and ended 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 (i.e., face value) of other tickets in the same quality tier as the auctioned tickets.¹²

¹²Auction bids are inclusive of convenience fees, whereas fixed-price tickets have separately stated face values and convenience fees. The starting bid in the auction is typically set equal to the face value plus convenience fees of tickets in the same quality tier as the auctioned tickets, 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 quality tier as the auction was \$225 and convenience fees were \$21.60, for an all-in fixed price of \$246.60, and the auction starting bid was set to be \$250. Throughout the paper, when we refer to a ticket’s face value or fixed price, we mean the price inclusive of convenience fees.

Figure 2: User interface for the Ticketmaster primary-market auction



Notes: This screenshot was taken from the Ticketmaster primary-market auction for tickets to see The Police in Fenway Park, Boston, July 29, 2007.

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 had just increased from \$530 to \$540, any bidders of \$530 would just have received a notification.

Starting in April 2007, bidders could also specify a quality threshold indicating the lowest quality group that their bids were valid for (e.g., valid only for the first 3 quality groups). As before, at the end of the auction, 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.

Relationship to Online Advertising 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. There are two central auction design differences between TM’s auction and the most well-known position auction format, the generalized second-price (GSP) auction studied by Edelman et al. [2007] and Varian [2007]. First, in TM’s auction successful bidders pay their own bid, rather than the next-highest bid. TM perceived that the benefit of using pay-your-own-bid is that it is simpler to explain to its customers. Second, bids are a per-ticket dollar amount, whereas typically in sponsored search auctions 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.

In Appendix B we modify the position auctions model of Edelman et al. [2007] and Varian [2007] to account for these two differences, and derive two basic results.¹³ First, the TM auction has a monotonic 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 we will discuss a weakness of the design below in Section 6, as well as some potential modifications. The results also can be interpreted as providing general support for the use of auctions over fixed prices, since fixed-price mechanisms satisfy none of these properties.

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.

¹³As in Edelman et al. [2007], we consider both sealed-bid and ascending versions of the model.

4.1 Primary-Market Auction Data

Our primary-market auction data cover all Ticketmaster auctions for concert tours that started in 2007. There are 22 concert tours, 576 concerts, and 759 auctions.^{14,15} 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 (22,348 in total). For each winning bid, we observe the following bid-level variables: customer identification number, bid amount, number of tickets (typically 2 or 4), time of bid, the section, row and seat numbers assigned to the bid, and the discrete quality group associated with the assigned tickets, per TM’s internal ranking. We also observe the following auction-level variables: artist, event date, event city, starting bid, ticket face value, starting time, and ending time. We use face values inclusive of all convenience fees, since bid amounts in the auction are inclusive of all fees. 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 quality tier as the auctioned tickets.

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 22 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 separate Perl script to extract several kinds of data from each html file.

Our focus is on successful eBay listings. For each such listing we observe concert-level data, specifically artist, event date and city, and data on the precise tickets being auctioned, specifically section, row, and number of tickets.¹⁶ We also observe eBay selling format parameters, such as

¹⁴For the majority of 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 TM’s original bidding language (pre April 2007), such a difference would have caused there frequently to be bidders with negative realized surplus.

¹⁵We drop all concerts in Canada. These concerts comprise less than 2% of our primary-market dataset.

¹⁶For both concert-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.

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

Winning bid: US \$999.99

Ended: Mar-21-07 14:41:40 PDT

Shipping costs: US \$20.00
Other (see description)
Service to [United States](#)

Ships to: United States

Item location: Boston, MA, United States

History: [1 bid](#)

Winning bidder: [zeis23 \(1042\)](#) ★

You can also: [Email to a friend](#) | [Sell one like this](#)

Listing and payment details: [Show](#)

Meet the seller
Seller: [eg201 \(1426\)](#) ★ **Power Seller**
Feedback: 100% Positive
Member: since Jun-08-99 in United States
■ [Read feedback comments](#)
■ [Ask seller a question](#)
■ [Add to Favorite Sellers](#)
■ [View seller's other items](#)

Buy safely
1. Check the seller's reputation
Score: 1426 | 100% Positive
[Read feedback comments](#)
2. Learn how you are protected
PayPal Free PayPal Buyer Protection.
[See eligibility](#)

Description (revised)

Item Specifics - Tickets

Event Name: **The Police** Number of Tickets: 2
Event Type: **Concerts**
Rock/Pop/Altern.

Venue State/Province: **Massachusetts**
Venue City: **Boston**
Venue Name: **Fenway Park**
[Click to View Map](#)

Month: **July**
Day: **29**
Year: **2007**
Time: **TBA**

Season Tickets: --

Stock Photo

Portions of this page provided by TicketTechnology. All rights reserved.

You are bidding on 2 SOLD OUT tickets to see the Police at Fenway Park in Boston on Sunday July 29th. What better way to see the show from 3rd Row Center on the Floor?? These are HARD TICKETS in hand ready to be shipped. Shipping will be twenty dollars via overnight. Any questions feel free to ask.

the auction's opening bid and/or Buy-it-Now price,¹⁷ and the listing's total selling price. We then reduce this total selling price by eBay's transaction fees, which were roughly 4% at the time of our data,¹⁸ and then divide by the number of tickets to obtain a per-ticket net-of-fees selling price.

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

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

¹⁷A Buy-it-Now price is a price which, if bid, ends the auction immediately. Buy-it-Now can be used by sellers to run a pure fixed-price listing (e.g., set the Buy-it-Now price equal to \$100, and set the auction's opening bid equal to \$100 as well) or to run a hybrid auction / fixed-price sale. For more on Buy-it-Now prices, see Budish and Takeyama [2001] and Milgrom [2004]. For additional details on eBay rules and on the use of eBay data in economic research, see Bajari and Hortagsu [2003, 2004].

¹⁸At 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 whether the winning bidder used PayPal to pay the seller; 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 about 4% of the \$580 transaction value.

that this particular seller did fill in), but that there is no such data field for seat number.¹⁹ 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, because prices for individual tickets are not representative of per-ticket prices for sets of two or more tickets (most consumers wish to attend concerts in groups rather than alone). We also drop observations in which the seller elected to use eBay’s “Dutch auction” format for selling variable quantities of tickets, because eBay sellers are inconsistent about whether they complete the “Number of Tickets” field on the eBay webpage based on the number of tickets awarded per winning bid (typically 1,2 or 4) or based on the total number of tickets the seller has available. As a result we are unable to reliably compute the price paid per ticket. Together, single-ticket and Dutch auction observations comprise about 14% 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 was standard practice in the secondary market for tickets at the time of our data, 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”).

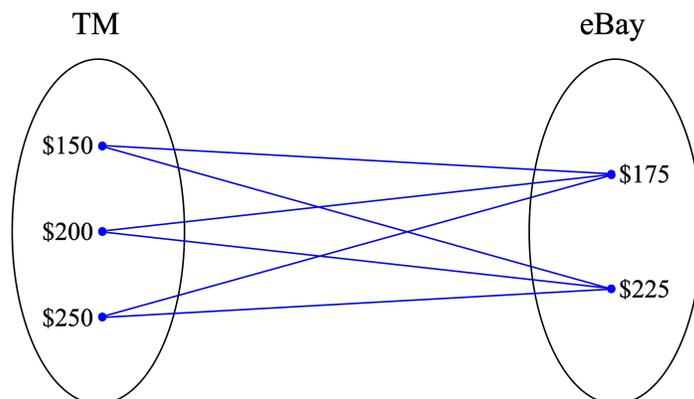
Second, eBay section and row data are input by eBay sellers, and are non-standardized. For instance, eBay sellers typically use the string “1” in the Row field to describe tickets that are in

¹⁹Listing 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.

²⁰In addition to typically being of negligible importance, seat data are 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.

²¹While 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. As a robustness exercise, we omitted these venues from the analysis; the results moved very little.

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



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 row input strings into standardized terms; e.g., all of the Row entries listed above get translated into “1”. We then create a venue-specific section dictionary, which translates each observed eBay section input string into a section name that appears on the seating chart of the venue for the event in question. Last, we match the eBay data to the Ticketmaster data at the level of concert-section-row, using the two dictionaries.

Third, when a particular c-s-r 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, with primary-market prices of \$150, \$200 and \$250, and a common secondary-market value of \$200. For robustness, we also consider a specification that instead aggregates the TM

Table 1: Summary statistics for matched primary- and secondary-market data

	Full TM Data Set	Matched Data Set	% Matched	
Concerts	576	464	80.56%	
c-s-r tuples	5,796	1,645	28.38%	
TM transactions	22,348	8,425	37.70%	
eBay transactions	N/A	3,532	N/A	
Matched Data Set	Mean	Std Dev	25th Perc	75th Perc
TM transactions per c-s-r tuple	5.12	7.45	2	6
eBay transactions per c-s-r tuple	2.15	5.92	1	2

Notes: “c-s-r” stands for concert-section-row. For details, see Sections 4.1-4.3.

transactions at the c-s-r level and treats each eBay observation separately, a specification that aggregates both TM and eBay transactions at the c-s-r level, and a specification that compares the minimum TM auction price in a c-s-r to the average eBay price in a c-s-r (cf. Appendix A.1). 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.

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

5.1 No Arbitrage

Figure 1 (in the introduction) presents a scatterplot of our matched dataset at the level of concert-section-row. In panel (a), the horizontal axis denotes the average price per ticket in the TM primary-market auction for the c-s-r, and the vertical axis denotes the average price per ticket in the eBay secondary market for the c-s-r. The vertical distance between a point and the 45-degree line represents the average profits associated with resale for that c-s-r. Panel (b) is identical, except that the horizontal axis denotes the 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.

The reasonably close fit of the data in panel (a) to the 45-degree line – especially in contrast to the data in panel (b) – conveys both that primary-market auction prices are informative of secondary-market prices and that average resale profits are small. Figure 5 presents a histogram of these resale profits. The mean resale profit is \$6.07, or 2.2% of the mean primary-market auction price of \$274.35. The 95% confidence interval of this estimate, clustering errors at the concert level, is [-\$7.57, \$18.59];²² unclustered, the confidence interval is [\$2.93, \$9.20]. 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 Appendices A.1 and A.2 suggest that, if anything, the \$6.07 estimate is too high and mean arbitrage profits are slightly negative.²³

There are two other interesting features of the distribution of resale profits to highlight. First, there is substantial variance: while the mean arbitrage profits are close to zero, there are specific tickets where the secondary-market value turns out to be substantially higher than the primary-market auction price, and vice versa. This variance is of course consistent with no arbitrage, which is a statement about a speculator’s *expected* profits from participating in the TM auction. 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 will return to this asymmetry below in Sections 5.2 and 6.1.

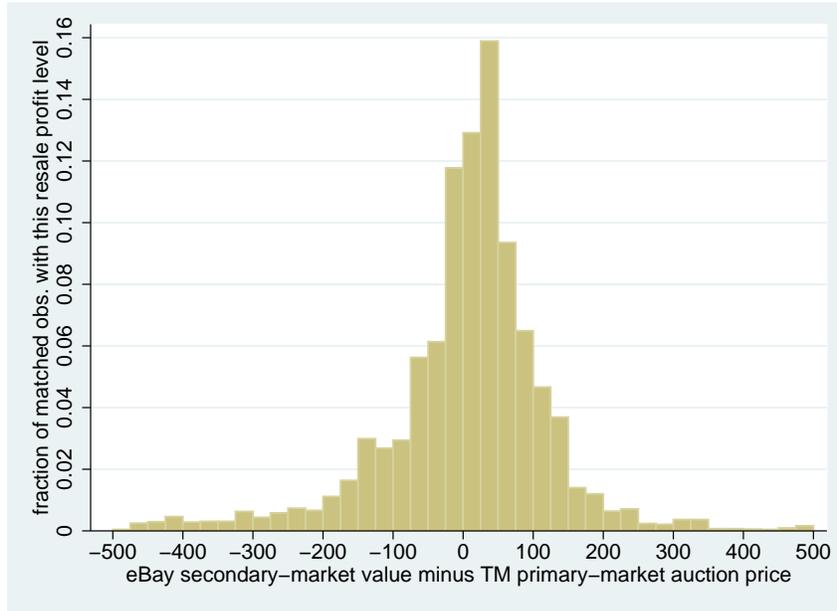
Table 2, column (1) regresses the eBay secondary-market value on the TM primary-market price. Interestingly, the regression best fit line is not the 45-degree line in Figure 1, panel (a). Rather, the regression best fit features a positive constant (\$47.38) and a slope less than one (0.85), with both differences statistically significant. This is another view of the same phenomena depicted in Figure 5, namely the positive mode and the fat left tail.

Table 2, column (2) regresses the eBay secondary-market value on the TM primary-market

²²The 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.

²³Over our four matching specifications, the 95% confidence intervals admit estimates of net arbitrage profits ranging from -\$38.52 to +\$18.59. If we assume that all sellers pay PayPal fees in addition to standard eBay fees, the 95% confidence interval for arbitrage profits becomes [-\$16.10, +\$10.86] under the main specification. For full details, see Appendices A.1 and A.2.

Figure 5: Distribution of resale profits



Notes: Resale profits are calculated as the difference between the eBay secondary-market value and the Ticketmaster primary-market auction price. Profits are on a per-ticket basis and are net of eBay fees. For more details, see the text.

Table 2: Price-informativeness regression results

eBay Secondary-Market Value	(1)	(2)	(3)
TM primary-market auction price	0.85 (0.07)		0.80 (0.09)
TM face value		1.74 (0.28)	0.33 (0.18)
Constant (\$)	47.38 (16.90)	28.31 (31.96)	13.57 (17.73)
R^2	0.66	0.24	0.66
Akaike Information Criterion	12.73	13.53	12.71
Schwarz Information Criterion	31,117	37,825	30,975

Notes: The dependent variable is eBay secondary-market value. For details, see the text.

face value. The interesting things to note are not the coefficients themselves but rather the informativeness measures. By all measures, face values are substantially less informative than auction prices. For instance, the R^2 in the face-value regression is just 0.24 as compared to 0.66 in the auction-price regression. Moreover, once we include primary-market auction prices in the regression of secondary-market value on primary-market price, there is very little additional information in ticket face values. By comparing columns (1) and (3) of Table 2, we see that inclusion of face values in the regression has little effect on measures of informativeness such as R^2 , the Akaike Information Criterion, or the Schwarz Information Criterion.

Face values are also systematically too low – this is the old and well-known underpricing phenomenon. This underpricing manifests most clearly in the scatterplot, with most of the mass in Figure 1, panel (b), being above the 45-degree line.²⁴ 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 \$8.5mm.

Altogether, our results confirm the basic benefits of using auctions over fixed prices: auction prices are more informative, raise more revenue, and nearly eliminate the arbitrage profits between the primary market and the secondary market.

5.2 Selection Concerns

Our eBay data are not a complete census of secondary-market activity, and for this reason one might worry that our results in Section 5.1 are affected by 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 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 in Section 5.1 suggest that arbitrage profits are small, 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

²⁴Observe that uninformativeness and underpricing are distinct phenomena. 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.

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 pure Buy-it-Now listings (+\$23.86) than for all other listings (-\$13.46).²⁵ Thus, our results do not appear to be driven by eBay sellers who set their Buy-it-Now prices too low.

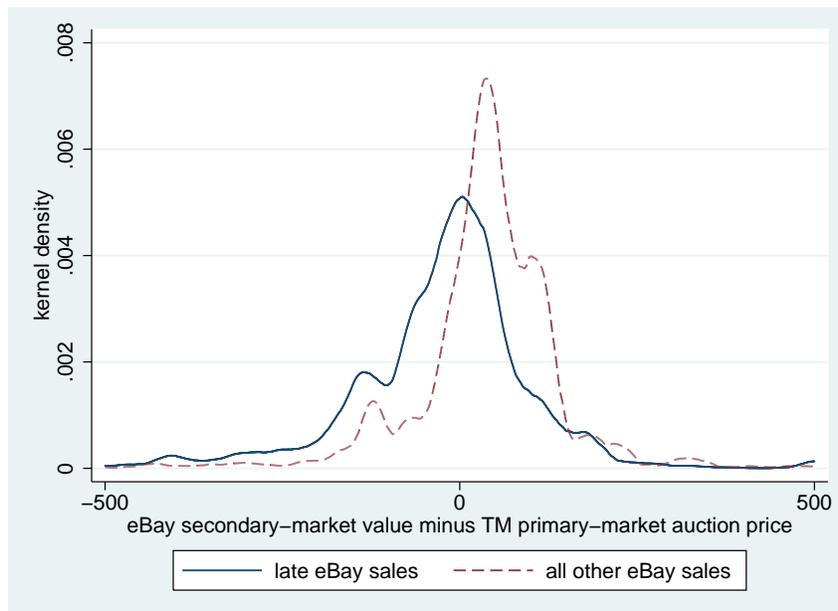
A third potential concern relates to our only seeing the portion of the aftermarket that occurs on eBay, as opposed to other venues such as StubHub. If the eBay component of the aftermarket is a random sample of the total aftermarket, then our use of just eBay data will cause us to have less power than if we had the full aftermarket, but will not cause bias. If the eBay component is non-random, however, this could cause bias. No arbitrage logic partly mitigates this concern: if eBay prices are systematically lower than StubHub prices (each net of fees), then arbitrageurs can buy on eBay to resell on StubHub (cf. Sweeting [2012]). However, we worry about the following: what if 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.²⁶ In this case eBay prices will be lower than average, which would cause us to under-estimate the returns to speculation. Given the direction of our results in Section 5.1 this is an important concern.

To address this third concern, we compare the distribution of arbitrage profits associated with resales on eBay that occur close to the event date, when we should worry about salvage-market effects, with the arbitrage profits associated with all other eBay resales. See Figure 6. eBay sales in the last 30 days before the event are associated with mean arbitrage losses of -\$24.16 per ticket, whereas the arbitrage profits associated with all other eBay sales are +\$40.93 per ticket. Notice as well that the early distribution has a higher mode than does the late distribution (small positive

²⁵This finding is consistent with results in Einav et al. [2013], which suggest that Buy-it-Now transaction prices are consistently higher than non-BIN transaction prices, across a wide variety of eBay categories.

²⁶The institutional reason why one might worry that sellers initially post on StubHub and then salvage on eBay, as opposed to the other way around, is the difference between 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. Pushing in the other direction, StubHub provides stronger buyer protection and transaction support than does eBay (e.g., ensuring that the buyer successfully receives the tickets from the seller), and these services may be especially valuable when the amount of time before an event is limited.

Figure 6: Resale profits, late eBay sales versus all other eBay sales



Notes: Late eBay sales are defined as eBay sales that occur within the last 30 days before the event. For more details, see the text.

profits), and that the late distribution has a fatter left tail (large losses). While these findings are consistent with the declining-price phenomenon documented by Sweeting [2012], they also suggest that we should be worried about salvage-market effects being present in our data.²⁷

A conservative response is to discard secondary-market data from the final days before the concert occurs as possibly tainted, and interpret the +\$40.93 mean arbitrage profits prior to the last 30 days as a conservative upper bound on arbitrage profits.²⁸ A second response is that failing to resell early is a real risk in this market (cf. Board and Skrzypacz [2010]),²⁹ and that the returns to speculation should be calculated based on the full sample of early and late eBay resales. A piece of evidence in support of this latter interpretation comes from looking at the difference between

²⁷Another 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 early eBay prices on TM primary-market auction prices is 0.77, versus 0.55 for late, and 0.66 for the full sample. This is consistent with results in Sweeting [2012], which show that the variance of secondary-market prices is much higher in the final days before an event.

²⁸The 95% confidence interval of this pre-last-30-days estimate is [\$26.64, \$54.28]. If we discard the final 15 days, rather than 30, the estimate is \$34.82 (95% CI: [\$21.63, \$47.87]).

²⁹Board and Skrzypacz [2010] characterize the optimal dynamic mechanism for sellers of perishable goods when buyers are forward looking. The optimal mechanism involves declining posted prices, followed by an auction in the final period. The auction can be interpreted as a salvage market, since its purpose is to ensure sale (modulo an optimally set reserve price) in the last period before the good expires. See Sweeting [2012] for further discussion of this paper and related dynamic pricing literature.

early eBay resales conducted using eBay’s pure fixed-price selling format and eBay’s pure auction format. Early pure BIN listings are associated with large, statistically significant profits of \$48.24 per ticket. Early pure auctions, by contrast, are associated with negative profits of -\$20.00 per ticket.³⁰ 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 [Board and Skrzypacz, 2010]. 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.³¹

6 Experienced vs. Inexperienced Bidders

6.1 Resale Profits

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 auctions, and then resells in the secondary market, he earns negligible profits.

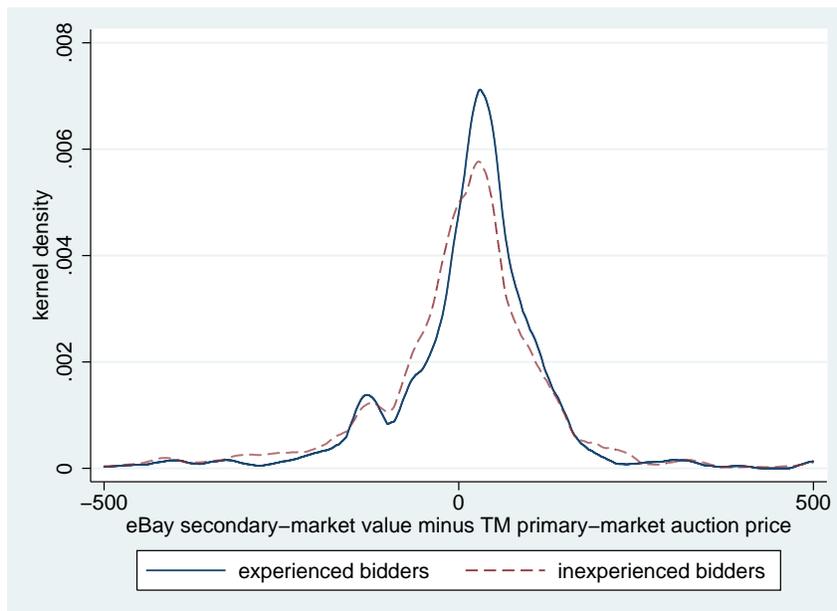
However, experienced bidders may have specialized knowledge about which tickets to purchase. Hence, to fully assess whether the auctions eliminate the rents of Bob the Broker, we need to look separately at the arbitrage profits of experienced and inexperienced bidders. If arbitrage profits are small on average, but large for professional resellers, this would cast the results of the previous section in a different light.

We exploit the fact that our TM data contain 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

³⁰Following Einav et al. [2013], 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 -\$20.00 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 -\$15.75.

³¹The early pure auctions exercise can also be interpreted as a response to the first selection concern described above. Since pure auctions nearly always result in a sale, one need not worry about bias from the use of only successful eBay listings. Under this interpretation, the -\$20.00 can be viewed as a conservative lower bound on arbitrage profits.

Figure 7: Resale profits, experienced versus inexperienced TM auction participants



Notes: Experienced TM auction participants are defined as participants who win tickets in at least 10 TM auctions. For more details, see the text.

the matched data). Such bidders account for 1% of the bidders in the TM data and roughly 16% of the transaction volume. We classify bidders who win between 1-9 auctions as inexperienced.³²

Figure 7 compares the distribution of resale profits for experienced and inexperienced bidders. While the distributions are remarkably similar overall, notice that the distribution for experienced bidders is to the right of the distribution for inexperienced bidders. The difference in means is statistically significant at the 1% level, with experienced bidders purchasing tickets with resale profits of \$19.49 per ticket, while inexperienced bidders purchase tickets with resale profits of \$2.47 per ticket.³³

The figure also suggests that experience accounts for some of the asymmetry in the distribution

³²Over half of the volume in the TM data is accounted for by bidders who win just a single TM auction (overall, not just restricted to matched data). The remaining 23% of volume corresponds to bidders with 2-9 transactions. We also consider a definition of experience based on the bidder winning at least 2 TM auctions in at least 2 cities (overall, not just restricted to the matched data). Such bidders account for 5% of bidders in the TM data and 24% of transaction volume. The results are very similar to our main specification. Last, we consider versions of both the 10+ auctions measure and the 2 auctions-2 cities measure based on bids rather than wins. Again, the results move very little. See Appendix A.3.

³³We performed a decomposition of the \$17.02 difference in profits between experienced and inexperienced bidders, and found that the difference is driven mostly by section-row selection within a concert (\$9.22). There are also differences from artist selection (\$1.68), concert selection (\$3.45), and paying a lower price than inexperienced bidders for seats in the same section-row (\$2.68).

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: 53.0% of transactions versus 42.4% of transactions (significant at 1%); and (ii) the most experienced bidders are significantly less likely to have large losses that exceed -\$100 per ticket: 11.7% of transactions versus 14.7% of transactions (significant at 5%). 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 perhaps 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 \$19.49 per ticket is large or small relative to time and effort costs, we emphasize that it is an order of magnitude smaller than the \$135.85 per ticket that bidders earn from resale under the counterfactual of using face values instead of the auction.

6.2 Overbidding

In addition to looking at the matched data, as above, we can also directly examine the TM bidding data for differences in bidding behavior between experienced and inexperienced bidders. In particular, given the pay-as-bid nature of the TM auction design and the criticism of such auctions in Friedman [1991] and Edelman and Ostrovsky [2007], we examine what we call “overbidding” – paying substantially more than is necessary to win tickets of a particular quality level.³⁴ We find evidence of occasional severe overbidding for tickets in the highest quality group in a given auction (e.g., 1st row): 13.91% of winning bids for tickets in the best quality group are at least 25% higher than was necessary to win seats in that group, 5.22% are at least 50% higher than was necessary, and 1.01% of winning bids are at least 100% higher than was necessary.³⁵ Table 3 shows that these overbids are rarely submitted by experienced bidders, and are disproportionately submitted

³⁴As can be seen in Figure 2, TM organizes tickets into quality groups in a manner that keeps within-group heterogeneity small. In a few auctions in our data, however, the heterogeneity of tickets within the highest quality group seemed to be non trivial.

³⁵Since there are many tickets per quality group, it is typically not possible to pay substantially more than other bidders for any quality group other than the first. For instance, in the auction depicted in Figure 2, a bid of \$1000 would win tickets in the highest quality group and pay substantially more than was necessary (\$540 as of this screen shot), whereas any bid between \$310 to \$420 would pay an amount within \$20 of the amount necessary to win tickets in the assigned row. In our data overall, 85.11% of winning bids are within \$0-\$10 of the next winning bid, and 94.09% are within \$0-\$50.

Table 3: Overbidding analysis

Cutoff Overbid Percentage	25%	50%	100%	Share in Full Dataset
% overbids at least this large	13.91%	5.22%	1.01%	($N = 6, 125$)
Of the overbids				
Experienced bidders	5.99%***	2.19%***	0.00%***	14.71%
Inexperienced bidders	94.01%***	97.81%***	100.00%***	85.29%
First auction	74.53%***	77.50%***	75.81%*	66.76%
Last auction	80.05%***	83.44%***	80.65%**	69.04%
Only auction	65.02%***	68.44%***	62.90%	56.05%

Notes: Overbidding summary statistics for the winning bids in the highest quality group in the TM primary-market auction. Each winning bid’s overbid percentage is calculated as (bid amount) / (lowest bid amount that won tickets in the highest quality group) - 1. The winning bid is included in the 25% (respectively, 50%, 100%) column if the overbid percentage is at least 25% (respectively, 50%, 100%). A bidder is defined as Experienced if he wins tickets in at least 10 TM auctions. First auction and Last auction are computed based on the date on which the auction ends. Only equals both First and Last. The column Share in Full Dataset refers to all winning bids in the highest quality group in the TM primary-market auction, not just overbids. “***” (respectively, “**”, “*”) means that the one-sided p -value of the difference between the figure reported in the overbidding column and the share in the full dataset is significant at the 1% (respectively, 5%, 10%) confidence level, based on a Bernoulli test. For more details, see the text.

by inexperienced bidders, especially for the overbids of at least 50%. We also find that overbidders disproportionately exit the market (i.e., they never bid again in another auction), though it is difficult to assign causality to this relationship.

6.3 Implications for TM’s Auction Design

It is not obvious whether TM and its clients should be concerned about overbidding by inexperienced bidders. On the one hand, overbidding by definition raises the artist’s revenue in a particular auction. On the other hand, overbidding is correlated with an effect – bidder exit – that is negative for the long-run health of the TM marketplace. Additionally, the risk of overbidding might deter some potential bidders from entering the TM auction market in the first place. This is analogous to the concern that Milton Friedman raised with respect to pay-as-bid US Treasury auctions, and which motivated Friedman’s proposal of uniform-price auctions as an alternative. In a uniform-price auction, Friedman wrote, “no one is deterred from bidding by fear of being stuck with an excessively high price” [Friedman, 1991]. Experience from other market design contexts also suggests that there are important benefits from reducing the strategic complexity of participating in a market [Roth, 2008]. For these reasons, we think it is constructive to propose alternatives to TM’s auction design

that mitigate its strategic complexity.

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 (GFP) to generalized second price (GSP). Specifically, if there are q_1, q_2, \dots, q_J tickets in each of J quality groups, the q_1 highest bidders receive tickets in the 1st quality group and pay the $q_1 + 1$ st highest bid, the next q_2 highest bidders receive tickets in the 2nd quality group and pay the $q_1 + q_2 + 1$ st highest bid, etc. Within each quality group, winning bidders are allocated tickets randomly. This is exactly like GSP except that there are multiple identical units available at each quality level. Hence, an appropriate name might be the generalized uniform price auction (GUP).³⁶

While the GUP protects bidders against severe overbidding, it is not strategy-proof (even in the large-market sense of Azevedo and Budish, 2013), for the same reason that the GSP is not strategy-proof (Edelman et al., 2007). A second possibility, then, would be to move to a fully strategy-proof 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 her demand – then a proxy version of the simultaneous ascending auction is strategy-proof in this context.³⁷ A disadvantage of the simultaneous ascending auction, relative to the GUP, is that it requires more preference information. Specifically, the auction would need to elicit each bidder’s value for each quality group, or block of quality groups. Given this information, the simultaneous ascending auction works by having each bidder’s proxy bid on the quality group that, at current prices, maximizes her surplus. For instance, if a bidder’s value is \$1000 for the 1st quality group and \$500 for the 2nd quality group, and current prices are \$800 and \$400 respectively, the proxy would bid on the 1st quality group, since $\$1000 - \$800 = \$200$ is more surplus than $\$500 - \$400 = \$100$. Specifically, the proxy would bid \$800 plus a single bidding increment. If the price of the 1st quality group increases to, say, \$950, the bidder’s proxy would

³⁶Recall that under the current TM auction design, higher bidders receive higher-quality tickets within the row. This convention might cause instability in a GUP, since any bidder who is not the highest bidder within her quality group has strict incentive to bid more: she receives strictly higher quality but pays the same amount. For this reason, and since quality heterogeneity within a row is typically quite small, we recommend randomizing the assignment within a quality group. This makes the prizes within a quality group homogeneous.

³⁷Without the single-mindedness assumption, the simultaneous ascending auction is not strategy-proof, but it does satisfy the large-market notion of approximate strategy-proofness defined in Azevedo and Budish [2013]. This follows from Proposition 1 of Azevedo and Budish [2013]. Intuitively, as the market grows 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. 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.

instead bid on the 2nd quality group, bidding \$400 plus a single bidding increment. Notice that this process involves the proxy sometimes lowering its bid amount in dollar terms, if it decides in response to new prices to target a lower quality level. 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 roughly double versus the fixed-price counterfactual; and, perhaps most importantly, the auctions eliminate or at least substantially reduce resale profits for speculators.

We conclude by speculating as to why, given that event-ticket underpricing and the associated rent-seeking behavior traces back to at least the 19th century, that it was only in the early 21st century that auctions were introduced to this market. While it is hard to know for sure, we speculate that the cause was a confluence of market design technology shocks associated with the rise of the internet.

Suppose that some artist genuinely wishes to underprice her tickets – set a price at which demand strictly exceeds supply. For instance, she is altruistic towards her fans and cares directly about their surplus (as in Che et al., 2013) or views fan concert surplus as complementary to long-run profit maximization (as in Mortimer et al., 2012). The ability to set a meaningful below-market price for “true fans” is undermined by primary-market and secondary-market technology shocks associated with the internet. In the primary market, the shift from physical and phone-based queues to online sales has enabled speculators to obtain a much larger share of the underpriced tickets, and hence of the surplus. As mentioned in the introduction, current estimates are that 20% of all primary-market tickets are obtained by speculators, with the figure as high as 90% in extreme cases. The increased accessibility of the secondary market, meanwhile, means that those fans who do obtain underpriced tickets in the primary market are likely to perceive the price in the easy-to-access online secondary market as the relevant opportunity cost, and hence may be

tempted to resell.

In addition, the internet was a shock to auction-design technology itself. It is difficult to imagine TM's position auction being run offline at a venue such as Sotheby's or Christie's, but it is straightforward and inexpensive to run online. Our story is thus: shocks to primary-market and secondary-market technology undermined the artist's ability to set a meaningful below-market price, and advances in auction-design technology enabled artists to accurately discover market-clearing prices.

Some indirect supporting evidence for this story is provided by examining two other nascent technologies in the event-tickets market, each of which can be interpreted as a substitute response to the same shocks. One alternative response is to continue to use posted prices, but use modern demand estimation technology to set these prices with greater accuracy. For instance, major league baseball teams host 80+ home games per year, and can use historical secondary-market data, historical sales patterns, and even knowledge about which pitchers are popular with fans to set prices for a particular game. Of course, this approach is not feasible for events that are relatively unique – e.g., Charles Dickens' reading tour – but for events where demand estimation is feasible, fixed prices can approximate the price discovery, revenue, and no-arbitrage performance of an auction. In addition, posted prices avoid the strategic complexity costs associated with TM's specific auction design.

A second alternative response is simply to deactivate the secondary market. We mentioned the 2007 Miley Cyrus / Hannah Montana tour in the introduction, in which tickets had a low face value, sold out in minutes, and then appeared on secondary-market sites at much higher prices, eliciting outrage from disappointed pre-teens, their parents, and several state Attorneys General. For the artist's³⁸ next tour, in 2009, Disney again set below-market prices, but this time adopted technology that eliminated 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 tickets were non-transferable because they were attached to a specific credit card, that had to be presented in

³⁸A point of clarification: in 2007 Miley Cyrus toured as both herself and her fictional alter ego, Hannah Montana. This was called the "Best of Both Worlds Tour", for obvious reasons. In 2009 Miley Cyrus toured only as herself, sans alter ego.

person at the concert venue.^{39,40} Turning off the secondary market is also a substitute for auctions: in either case, there are no excess rents for Bob the Broker.

³⁹See Waddell [2009] regarding Miley Cyrus / Hannah Montana's use of this technology, called "paperless ticketing", and see Lipka [2011] regarding lobbying efforts by an entity called the Fan Freedom Project (backed in part by eBay and StubHub) to make paperless ticketing illegal.

⁴⁰The 2012 Summer Olympics in London provide a cautionary tale regarding the use of paperless tickets (see Economist, 2012). Tickets in the primary market were allocated in large part to corporate sponsors, who frequently discover at the last minute that they are unable to attend. As a result, there were large blocks of empty seats at the Olympics, which was both wasteful and embarrassing for the event's organizers. An interesting question for future research is how best to design such a ticketing system; presumably, optimal design incentivizes ticket-holders who are unable to attend the event to return their tickets back to the center, but in a way that does not induce speculative behavior. Airlines accomplish this goal using a combination of refunds and cancellation fees; a difference between event tickets and airline tickets is that airlines do not care per se about filling the plane, only about revenue, whereas event organizers may care per se about filling the venue.

References

- Lawrence Ausubel and Paul Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):1–42, 2002.
- Eduardo Azevedo and Eric Budish. Strategy-proofness in the large. Working Paper, 2013.
- 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.
- 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.
- Eric Budish and Lisa Takeyama. Buy prices in online auctions: Irrationality on the internet? *Economics Letters*, 72:325–33, 2001.
- Jeremy Bulow and Paul Klemperer. The generalized war of attrition. *American Economic Review*, 89(1):175–189, March 1999.
- Jeremy Bulow and Paul Klemperer. Regulated prices, rent-seeking, and consumer surplus. *Journal of Political Economy*, 120(1):160–186, February 2012.
- Yeon-Koo Che, Ian Gale, and Jinwoo Kim. Assigning resources to budget-constrained agents. *Review of Economic Studies*, 80(1):73–107, July 2013.
- 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*, 66(1):167–192, 2000.
- Pascal Courty. Some economics of ticket resale. *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, August 4, 2012. URL <http://www.economist.com/node/21559937>.
- 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. Working Paper, February 2013.
- Milton Friedman. How to sell government securities, August 28, 1991. URL <http://0055d26.netsolhost.com/friedman/pdfs/wsaj/WSJ.08.28.1991.pdf>.
- Nathan Hubbard. An open letter to fans and partners on ticket reselling, August 2011. URL <http://blog.ticketmaster.com/2011/08/>.
- Live Nation Entertainment Inc. Fiscal year 2012 Form 10-K, February 2013.
- 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. Resale and rent-seeking: An application to ticket markets. *Review of Economic Studies*, Forthcoming, 2013.

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. Fiscal year 2009 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. Technical report, Forrester Research, 2008.

Chris Nelson. Ticketmaster auction will let highest bidder set concert prices, September 1, 2003. URL <http://www.nytimes.com/2003/09/01/business/media-ticketmaster-auction-will-let-highest-bidder-set-concert-prices.html?pagewanted=all>.

Joe Nocera. Rigging the i.p.o. game, March 2013. URL <http://www.nytimes.com/2013/03/10/opinion/sunday/nocera-rigging-the-ipo-game.html>.

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.

Ellen Rosen. In the race to buy concert tickets, fans keep losing, October 6, 2007. URL <http://www.nytimes.com/2007/10/06/business/06money.html>.

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, September 12, 2006. URL <http://online.wsj.com/article/SB115802470458860177.html>.

Steve Stecklow. Concerts: Can't get no ... tickets?, January 7, 2006. URL <http://online.wsj.com/article/SB113658907268840218.html>.

Andrew Sweeting. Price dynamics in perishable goods markets: The case of secondary markets for major league baseball tickets. *Journal of Political Economy*, 120(6):1133–1172, December 2012.

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

New York Times. Mr. Dickens in Boston - the excess demand for tickets, November 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-idUSTRE55KOCX20090621>.

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

Appendix

A Robustness Tests

A.1 Robustness of Main Results to Alternative Matching Specifications

As discussed in Section 4.3, there are four potential ways to match primary- and secondary-market observations within the same concert-section-row tuple. In our main specification, we match each TM primary-market transaction with the average price of the eBay secondary-market transactions for the c-s-r in question. This approach allows us to exploit all of the variation in the winning bids in Ticketmaster’s high-quality dataset. In this section, we consider the robustness of our results to three alternative matching specifications. First, we match each secondary-market transaction with the average price of the primary-market transactions for the c-s-r in question. In the context of Figure 4, the hypothetical c-s-r tuple would represent two observations, with secondary-market prices of \$175 and \$225, and a common primary-market price of \$200. Next, for each c-s-r, we match the average price of the primary-market transactions with the average price of the secondary-market transactions. Therefore the hypothetical c-s-r in Figure 4 would represent one observation, with primary- and secondary-market prices of \$200. Last, for each c-s-r, we match the minimum primary-market auction price with the average secondary-market price. We consider this specification to address the possibility that, within a c-s-r, bidders who paid higher prices in the primary-market auctions were less likely to resell their tickets in the secondary market. In this case, the hypothetical c-s-r in Figure 4 would represent one observation, with primary- and secondary-market prices of \$150 and \$200, respectively.

Table A1 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 four above-discussed matching specifications. 95% confidence intervals are calculated using the bootstrap, with the data clustered at the concert level (cf. footnote 22).

Over all of our specifications, the 95% confidence intervals admit estimates of arbitrage profits, net of eBay transaction fees, ranging from -\$38.52 to +\$18.59.

Table A1: Summary statistics for arbitrage profits, and price discovery regression results, under different approaches to data aggregation.

Unit of Analysis	N	Avg. Profit (\$)	Clustered CI (\$)	Mode (\$)	Skewness	Price-Discovery Regression		
						Slope	Constant (\$)	R^2
TM Transaction (main specification)	8,425	6.07	[-7.57, 18.59]	28.46	-0.69	0.85 (0.07)	47.38 (16.90)	0.66
eBay Transaction	3,532	-20.45	[-38.52, -3.44]	22.89	-0.61	0.76 (0.06)	40.03 (14.98)	0.57
CSR Combination	1,646	-18.17	[-26.91, -9.42]	22.50	-1.75	0.77 (0.07)	43.64 (15.82)	0.65
CSR Combination (min. TM price)	1,646	-7.41	[-16.27, 1.35]	25.00	-1.09	0.79 (0.07)	46.82 (15.92)	0.63

Notes: For details on the different approaches to data aggregation, see Section 4.3. Confidence intervals are calculated at the 95% level. Standard errors are clustered at the concert level and are calculated via bootstrap. Modes are calculated from kernel density estimates that apply the Epanechnikov kernel function to a grid of 1,500 points on the interval [-\$500, +\$500]. The price-discovery regression results pertain to the regression of eBay secondary-market value on TM primary-market auction price.

A.2 Robustness of Main Results to Assumptions on eBay Transaction Fees

While all eBay sellers pay eBay fees (i.e., final-value fees and insertion fees), we do not observe whether a given eBay seller transacted using PayPal, or some other method such as cash or check. Therefore, in our main specification, we subtract eBay fees from the eBay transaction price, but do not consider PayPal fees, which were roughly a bit less than 3% of the eBay sale price at the time of our data.⁴¹

Table A2 reports gross profits, profits net of eBay fees, and profits net of both eBay fees and PayPal fees. Over all of our these specifications, the 95% confidence intervals admit estimates of profits ranging from -\$16.10 to +\$30.31.

⁴¹See footnote 18 for further details on PayPal fees.

Table A2: Summary statistics for arbitrage profits, and price discovery regression results, under different assumptions on eBay transaction fees.

	N	Avg. Profit (\$)	Clustered CI (\$)	Mode (\$)	Skewness	Price-Discovery Regression Slope	Constant (\$)	R^2
Resale Profits Gross of eBay Fees	8,425	17.03	[2.81, 30.31]	37.20	-0.48	0.87 (0.07)	52.98 (17.05)	0.66
Resale Profits Net of eBay Fees (main specification)	8,425	6.07	[-7.57, 18.59]	28.46	-0.69	0.85 (0.07)	47.38 (16.90)	0.66
Resale Profits Net of eBay Fees and PayPal Fees	8,425	-1.94	[-16.10, 10.86]	24.23	-0.97	0.83 (0.07)	45.80 (16.44)	0.66

Notes: For details on eBay fees, see Section 4.2. For all other aspects of the table see the notes to Table A1.

A.3 Robustness of Results on Bidder Experience

In Section 6, we show that experienced bidders, defined as those who win at least 10 TM auctions, earn modest positive arbitrage profits from buying in the TM auctions and reselling on eBay. The profits of experienced bidders are significantly higher than those of inexperienced bidders, namely, those who win less than 10 TM auctions. Since our classification of experienced and inexperienced bidders is somewhat arbitrary, we consider the robustness of our results on bidder experience to an alternative classification.

In particular, we define an experienced bidder as one who wins TM auctions for concerts in at least two different cities, performed by at least two different artists. Therefore inexperienced bidders are those who win tickets to just one event, those who are avid followers of a given artist (who follow just one artist across several cities), and those who simply enjoy live musical events (i.e., those who attend several concerts in a given city). The reasoning behind this alternative specification is that bidders who are not professional resellers are likely to fall into one of the three aforementioned categories of inexperienced bidders. Table A3 shows that the results on bidder experience are qualitatively unchanged when we employ our alternative classification. In both the main and the alternative specifications, i) experienced bidders make small positive profits (significant at 1% and 10%, respectively), ii) inexperienced bidders make essentially zero arbitrage profits, and iii) the profits of experienced bidders are significantly larger than those of inexperienced bidders (significant at 1%). Our results also remain qualitatively unchanged if we define either measure of bidder experience based on the number of auctions the bidder participated in rather than the number of auctions the bidder won.

Table A3: Summary statistics for arbitrage profits, and price discovery regression results, for different definitions of bidder experience.

	N	Avg. Profit (\$)	95% CI (\$)	Mode (\$)	Skewness	Price-Discovery Regression Slope	Constant (\$)	R^2
Experienced bidders								
≥ 10 transactions (main specification)	1,785	19.49	[5.32, 33.04]	29.94	0.75	0.95 (0.09)	32.03 (19.14)	0.67
≥ 2 cities and artists	2,597	15.02	[-1.96, 29.87]	30.17	-0.50	0.90 (0.09)	41.88 (21.12)	0.68
Inexperienced bidders								
< 10 transactions (main specification)	6,640	2.47	[-12.26, 15.65]	27.83	-0.89	0.83 (0.07)	49.17 (17.18)	0.65
< 2 cities or artists	5,828	2.09	[-12.66, 15.28]	27.16	-0.79	0.82 (0.06)	50.45 (16.10)	0.64

Notes: For details on bidder experience measures, see Section 6. For all other aspects of the table see the notes to Table A1.

B Model (not for publication)

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, risk-neutral bidders.⁴² Initially, we think of a bidder as either a fan who intends to use the tickets herself, 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 to $\alpha_k v_i$: v_i is bidder i 's type, and α_k describes the quality of ticket k , with $\alpha_1 > \alpha_2 > \dots > \alpha_K$.⁴³ 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 in the main text). 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 $k = 1, \dots, K$. There

⁴²For expositional purposes, we will refer to a pair of tickets simply as one object, or one ticket.

⁴³In Edelman et al. [2007], α_k is the click-through rate of the k^{th} slot, and v_i is the i^{th} advertiser's private value per click.

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, the k^{th} -highest 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 $b_{(k)}$: her 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 her [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 $k = 2, \dots, K$, gets the k^{th} ticket, and pays the amount at which she herself dropped. The $n - K$ bidders who do not get a 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. (*Efficiency of Sealed-Bid Auction*) *There exists a symmetric monotonic Bayes-Nash equilibrium of the sealed-bid TM auction in which all bidders bid according to*

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

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 $T(v; \underline{v}, k)$ denote the amount of time bidder v would be willing to wait before dropping out, conditional on the event that none of the other active bidders drop out during this time. Additionally, define the hazard rate in the standard manner, $h(v) = \frac{f(v)}{1-F(v)}$.

Proposition 2. (*Efficiency of Ascending Auction*) *The unique symmetric perfect Bayesian equilibrium of the ascending auction is defined by,*

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

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, $\frac{\partial T(v;v,k)}{\partial v}$, with their marginal benefit from doing so, $(\alpha_k - \alpha_{k+1}) vkh(v)$. 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, and the lowest type gets zero surplus, we immediately have the following corollary:

Corollary 1. *(Revenue Performance) The sealed-bid and ascending Ticketmaster auctions are revenue equivalent to any other efficient auction design in which the lowest type gets zero surplus, 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. Fans’ types are drawn independently and identically from the continuously differentiable distribution $F_{fan}(\cdot)$ with support $[0, \bar{v}]$. Fix $\varepsilon > 0$ and $w \in [\varepsilon, \bar{v}]$. 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

$$F(x) = \beta F_{pro}(x) + (1 - \beta) F_{fan}(x). \quad (3)$$

$F(\cdot)$ has support on $[0, \bar{v}]$, and we assume that $F(\cdot)$ is continuously differentiable $\forall \beta \in [0, 1]$.⁴⁴ This method of constructing a population comprising professional resellers and fans allows us to use the symmetric bidding equilibria that we derived in Propositions 1 and 2.

Let $n_{pro} \equiv \beta n$ and $n_{fan} \equiv (1 - \beta) n$ denote, respectively, the expected number of professional resellers and fans in the auction. We then model entry by speculators by allowing n_{pro} to increase while n_{fan} remains constant. This approach to endogenizing entry allows for the following simple no arbitrage statement.

Proposition 3. *(No Arbitrage) Let $s(v; n_{pro}, n_{fan})$ denote the expected surplus, conditional on winning some ticket, of a professional reseller with value v . Under either the sealed-bid or ascending TM auction, for any $v \in [w - \varepsilon, w]$ and any finite n_{fan} , $\lim_{n_{pro} \rightarrow \infty} s(v; n_{pro}, n_{fan}) = 0$.*

Proposition 3 formalizes in a simple manner that free entry by speculators causes them to earn negligible resale profits, even when we condition on the event that they win a ticket in the auction. Though a simple result, it highlights an important difference between auctions and fixed-price selling mechanisms.

⁴⁴Since the support of F_{pro} is a subset of that of F_{fan} , this is equivalent to assuming that (i) $F'_{pro}(w - \varepsilon) = 0$ or $w = \varepsilon$, and (ii) $F'_{pro}(w) = 0$ or $w = \bar{v}$.

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:

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

where

$$P_k(x) = \binom{n-1}{k-1} F(x)^{n-k} (1-F(x))^{k-1} \quad (4)$$

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,

$$\sum_{k=1}^K ([v\alpha_k - b(v)] P'_k(v) - P_k(v) b'(v)) = 0. \quad (5)$$

We can use (5) to solve for the equilibrium bidding function as follows. Rearranging terms,

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

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

$$b(v) = \frac{1}{\sum_{k=1}^K P_k(v)} \int_0^v x \sum_{k=1}^K \alpha_k P'_k(x) dx. \quad (6)$$

We can then use integration by parts to derive (1):

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

To finish the proof we must confirm, as we had assumed above, that the bidding function (1) is strictly increasing. Using (6), notice that

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

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

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

Proof. Let $x \in (0, \bar{v})$. Define $\alpha_{K+1} \equiv 0$. We have

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

The first equality follows from the definition of $P'_k(\cdot)$, and the second results from algebraic manipulation. The third 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

$$\begin{aligned} & \left(\sum_{k=1}^K P_k(v) \right) \left(\sum_{k=1}^K \alpha_k P'_k(v) \right) - \left(\sum_{k=1}^K P'_k(v) \right) \left(\sum_{k=1}^K \alpha_k P_k(v) \right) \geq 0 \\ & \iff \sum_{j=1}^K \sum_{k=1}^K P_k(v) \alpha_j P'_j(v) - \sum_{j=1}^K \sum_{k=1}^K P'_j(v) \alpha_k P_k(v) \geq 0 \\ & \iff \sum_{j=1}^K \sum_{k=1}^K (\alpha_j - \alpha_k) P_k(v) P'_j(v) \geq 0 \\ & \iff \sum_{j=1}^{K-1} \sum_{k=j+1}^K \underbrace{(\alpha_j - \alpha_k)}_{>0, \text{ since } j < k} [P_k(v) P'_j(v) - P_j(v) P'_k(v)] \geq 0. \end{aligned}$$

where the second, third and fourth inequalities each follow from rearranging terms. It is thus sufficient to show that, for any $v \in [0, \bar{v}]$ and any $j < k$,

$$P_k(v) P'_j(v) - P_j(v) P'_k(v) \geq 0.$$

Note that

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

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

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

as required. Therefore, the bidding function (1) is indeed strictly increasing. This in turn implies that the resulting allocation is efficient. \square

Proof of Proposition 2

Proof. The proof will proceed by induction. Bidders draw their private valuations from $F(\cdot)$, and the mechanics of the auction are as described previously. 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 object, and pay the price at which the last drop-out occurred.

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. 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. In each such auction, the bidders are the winners of the previous auction, and an upgrade of one quality level is awarded to all bidders but one. The process naturally stops after the quality upgrade from object 2 to object 1 has been sold.

In aggregate, bidders who do not win any object pay nothing, and bidders who win an object 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:*

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

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 initial second-price English auction is characterized by:*

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

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 $T(v; \underline{v}, k) = v - \underline{v}$.

□

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), with $k = 1$.

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 the bidder in question knows that she will follow the proposed strategies in each subgame $j < k$, 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 $T(v; \cdot, \cdot)$ 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 . Then at the start of subgame $k - 1$, v is less than the lowest possible type of the other bidders who have not dropped out, \underline{v}_{k-1} . Since the bidder follows $T(v; \underline{v}_{k-1}, k - 1)$ 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 payoff 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 playing truthfully to bidding as type $v^* > v$. 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 $T(v; \cdot, k)$. 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); $\forall k \in \{1, 2, \dots, K\}$, $T(\cdot; \underline{v}, k)$ is clearly strictly increasing.

□

Proof of Proposition 3

Proof. Since the sealed-bid and ascending TM auctions are efficient, and the lowest type gets zero surplus in both auctions, Myerson's Lemma implies that all bidders have the same expected surplus in both auctions. Moreover, all bidders have the same expected surplus in both auctions conditional on losing (namely, zero). It follows that all bidders must also have the same expected surplus in both auctions conditional on winning some object. Hence any bidder's expected surplus conditional on winning some object, in either auction, can be expressed as

$$s(v; n_{pro}, n_{fan}) = \sum_{k=1}^K [v\alpha_k - b(v)] P_k(v; n_{pro}, n_{fan}),$$

where $b(\cdot)$ is the sealed-bid auction equilibrium bidding function. Using (1), we find that

$$s(v; n_{pro}, n_{fan}) = \sum_{k=1}^K \int_0^v \alpha_k P_k(x; n_{pro}, n_{fan}) dx,$$

$$\begin{aligned} \text{where } P_k(x; n_{pro}, n_{fan}) &= \binom{n_{pro} + n_{fan} - 1}{k-1} \left(\frac{n_{pro} F_{pro}(x) + n_{fan} F_{fan}(x)}{n_{pro} + n_{fan}} \right)^{n_{pro} + n_{fan} - k} \\ &\quad \times \left(1 - \frac{n_{pro} F_{pro}(x) + n_{fan} F_{fan}(x)}{n_{pro} + n_{fan}} \right)^{k-1}. \end{aligned}$$

The last expression uses (3) and (4), noting that (i) the total number of bidders is $n = n_{pro} + n_{fan}$ and (ii) the expected share of professional resellers is $\beta = \frac{n_{pro}}{n_{pro} + n_{fan}}$. Since $\left(1 - \frac{n_{pro} F_{pro}(x) + n_{fan} F_{fan}(x)}{n_{pro} + n_{fan}} \right)^{k-1} \leq 1$, $\forall k \in \{1, 2, \dots, K\}$,

$$P_k(x; n_{pro}, n_{fan}) \leq \binom{n_{pro} + n_{fan} - 1}{k-1} \left(\frac{n_{pro} F_{pro}(x) + n_{fan} F_{fan}(x)}{n_{pro} + n_{fan}} \right)^{n_{pro} + n_{fan} - k}.$$

Now define $G(x) \equiv \max\{F_{pro}(x), F_{fan}(x)\}$. Then $\forall x < w$,

$$\frac{n_{pro} F_{pro}(x) + n_{fan} F_{fan}(x)}{n_{pro} + n_{fan}} \leq G(x) < 1.$$

Hence, $\forall x < w$,

$$\begin{aligned} P_k(x; n_{pro}, n_{fan}) &\leq \binom{n_{pro} + n_{fan} - 1}{k-1} G(x)^{n_{pro} + n_{fan} - k} \\ &\leq \frac{(n_{pro} + n_{fan} - 1)^{k-1}}{(k-1)!} G(x)^{n_{pro} + n_{fan} - k} \\ &= \frac{(n_{pro} + n_{fan} - 1)^{k-1}}{(k-1)! G(x)^{k - (n_{pro} + n_{fan})}} \end{aligned}$$

The second inequality and the equality follow from simple algebraic manipulations. Since $G(x) < 1$, the last step converts the expression in question into $\frac{\infty}{\infty}$ form when $n_{pro} \rightarrow \infty$, allowing us to use l'Hôpital's Rule.

Applying this rule $k - 1$ times yields

$$\begin{aligned} \lim_{n_{pro} \rightarrow \infty} \frac{(n_{pro} + n_{fan} - 1)^{k-1}}{(k-1)!G(x)^{k-(n_{pro}+n_{fan})}} &= \lim_{n_{pro} \rightarrow \infty} \frac{G(x)^{n_{pro}+n_{fan}-k}}{(-\log[G(x)])^{k-1}} \\ &= 0. \end{aligned}$$

We can now conclude that, $\forall x < w$, and $k \in \{1, 2, \dots, K\}$, $\lim_{n_{pro} \rightarrow \infty} P_k(x; n_{pro}, n_{fan}) = 0$. Therefore $\forall v \in [w - \varepsilon, w)$,

$$\begin{aligned} \lim_{n_{pro} \rightarrow \infty} s(v; n_{pro}, n_{fan}) &= \sum_{k=1}^K \int_0^v \alpha_k \left[\lim_{n_{pro} \rightarrow \infty} P_k(x; n_{pro}, n_{fan}) \right] dx \\ &= 0. \end{aligned}$$

In order to complete the proof, we must address the case in which $v = w$. Note that

$$\begin{aligned} \sum_{k=1}^K \alpha_k P_k(w; n_{pro}, n_{fan}) &\leq \alpha_1 \sum_{k=1}^K P_k(w; n_{pro}, n_{fan}) \\ &\leq \alpha_1, \end{aligned}$$

because $\sum_{k=1}^K P_k(w; n_{pro}, n_{fan}) \leq 1$. Since α_1 is finite, and since $\lim_{n_{pro} \rightarrow \infty} P_k(x; n_{pro}, n_{fan}) = 0 \forall x < w$ and $k \in \{1, 2, \dots, K\}$, we have

$$\begin{aligned} \lim_{n_{pro} \rightarrow \infty} s(w; n_{pro}, n_{fan}) &= \sum_{k=1}^K \int_0^w \alpha_k \left[\lim_{n_{pro} \rightarrow \infty} P_k(x; n_{pro}, n_{fan}) \right] dx \\ &= 0 \end{aligned}$$

as well, as required. □