American Economic Journal: Microeconomics 2023, 15(1): 142–170
https://doi.org/10.1257/mic.20180230

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

By Eric Budish and Aditya Bhave*

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? This paper studies the introduction of auctions into this market by Ticketmaster in the mid-2000s. By combining primary-market auction data from Ticketmaster with secondary-market resale value data from eBay, we show that Ticketmaster's auctions "worked": they substantially improved price discovery, roughly doubled performer revenues, and, on average, nearly eliminated the potential arbitrage profits associated with underpriced tickets. We conclude by discussing why, nevertheless, the auctions failed to take off. (*JEL* D44, D47, L82)

> 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)

* Budish: University of Chicago Booth School of Business (email: eric.budish@chicagobooth.edu); Bhave: Bank of America Merrill Lynch (email: aditya.bhave.17@gmail.com). Michael Ostrovsky was coeditor for this article. Alan Krueger made this study possible, helping a graduate student he barely knew who shared his love of music and markets. This paper is dedicated to his memory. 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 the event-ticket market. We are grateful to Adam Juda and Aaron Roth for the Perl scripts we used to obtain eBay data. We thank Glenn Ellison and Steve Tadelis for valuable conference discussions. We thank two anonymous referees for valuable feedback. 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, and Hal Varian for helpful conversations about this research, 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, ACM EC 2013, the 2015 Econometric Society winter meetings, and the 2018 Music Industry Research Association annual conference. We thank Meru Bhanot, Mika Chance, Natalia Drozdoff, Sam Hwang, Willy Lin, Yi Sun, and Cameron Taylor for excellent research assistance. 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.

† Go to https://doi.org/10.1257/mic.20180230 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

> 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)

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* 1867a, b). 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 2014). For instance, when the Disney star Miley Cyrus (aka Hannah Montana) first toured the United States in 2007–2008, tickets with a face value of at most \$64 sold out in approximately 12 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). Tickets for the hit broadway show "Hamilton" sold out so quickly (sometimes in seconds) and were so expensive in the secondary market (Mankiw 2016) that its creator and star Lin-Manuel Miranda began publicly advocating for changes to state ticket broker laws (Miranda 2016).

Economists have long found this phenomenon puzzling (Becker 1991; Landsburg 1993; Rosen and Rosenfield 1997; Courty 2003a; Baliga 2011; Krueger 2019). First, underpricing reduces revenues for the performer. Second, underpricing encourages socially wasteful rent-seeking by ticket speculators. While it 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, Gale, and Kim 2013; Courty 2019; though see also Bulow and Klemperer 2012), or the sale of complementary goods over the artist's life cycle (Mortimer, Nosko, and Sorensen 2012)—it is hard to argue that artists genuinely wish for *ticket speculators* to get tickets at below-market prices and earn arbitrage profits. In that sense, the true puzzle is the *combination* of low prices and rent-seeking by speculators due to an active secondary market.

Modern information technology has only exacerbated the scale of this 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. In the Dickens era, and as recently as the late twentieth 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 State Office of the Attorney General 1999). In the present day, tickets can be amassed in the primary market online using software bots and sold in the secondary market on websites such as eBay and StubHub (Zetter 2010; Ticketmaster Blog 2011b). Industry estimates suggest that on the order of 20 percent of all tickets purchased in the primary market are resold in the secondary market, constituting on the order of $15 billion of volume annually. In extreme cases, speculators amass as many as 90 percent of the tickets available for a particular event.[1] In the colorful words of the Arkansas attorney general, "All hell broke loose with Hannah Montana" (Rosen 2007).

Economic theory suggests that there are two basic choices for how to curb this rent-seeking. One choice is to make tickets nontransferable or otherwise ban resale. This would allow artists to set whatever price they choose for their ticket, including an artificially low price (e.g., "Hamilton" awards 46 seats per night at $10 each to fans who win a lottery and prevents resale for these tickets (Paulson 2016)). We will return to this idea in the conclusion, and also point the interested reader to Courty (2019), but note here that the legality of banning resale is fiercely debated and that an entire lobbying organization, the Fan Freedom Project—funded in part by eBay and StubHub—is dedicated to preserving so-called "resale rights" (see Lipka 2014; Budish 2019).

The second choice, of course, is to set a market-clearing price. This paper studies an effort to do just that—Ticketmaster's 2003 introduction of primary-market auctions for concert tickets. Ticketmaster's chief executive officer remarked at the time:

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

This paper shows that Ticketmaster'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 Ticketmaster (abbreviated "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.[2] 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

[1] The 20 percent figure is from a blog post by Ticketmaster's CEO on August 12, 2011 (Ticketmaster Blog 2011a). The $15 billion figure is from Tan (2016), with similar figures reported elsewhere. The 90 percent figure is a Ticktmaster estimate reported to the authors, based on software used to detect the software bots mentioned in the text. To give a sense of the growth of secondary-market activity, Leslie and Sorensen (2014) report that the rate of resale was on the order of 5 percent in summer 2004 and report industry estimates that the overall secondary-market size was about $3 billion. For further context on these figures and additional institutional detail, see Budish (2019).

[2] For 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.

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.[3] Given these three types of data, it is straightforward to calculate the TM auctions' effect on price discovery, revenues, and arbitrage profits.

Figure 1 conveys our main results. The left panel is a scatterplot of TM primary-market auction prices and subsequent eBay secondary-market resale values; each dot is a concert-section-row tuple (e.g., The Police, 7/29/2007, Section A3, Row 2). 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 resale value. The mean difference between the auction price and the subsequent resale value is just $6.07, or around 2 percent 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 now $136, representing a 94 percent return on the average primary-market face value of $145.[4] 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. The auctions essentially eliminate the scope for rent-seeking, at least on average.

We also explore the auction performance of the most experienced bidders in the TM auction—the top 1 percent of bidders, who account for 16 percent 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 firms use for keyword advertising (Edelman, Ostrovsky, and Schwarz 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

[3] The reason that we match tickets at the level of section and row, rather than the precise seat within the row, is that for privacy reasons eBay listings do not indicate the specific seat. This was standard practice in the secondary market at the time of our data. The reader should interpret our eBay data as providing a measure of the resale value of tickets sold in TM's auction, not necessarily as actual resales of tickets purchased in the TM auction. See Section IVC for more details.

[4] The large magnitude is consistent with Leslie and Sorensen's (2014) finding that the most severe underpricing occurs for high-quality tickets, which are the focus of TM's auctions.

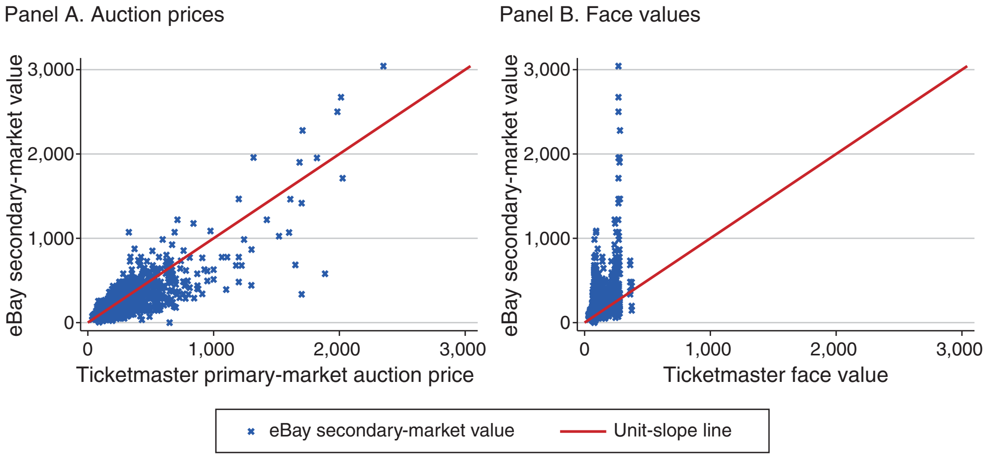

Figure 1. Relationship between eBay Secondary-Market Values and Ticketmaster Primary-Market Prices, Both Auction Prices and Face Values

Notes: Panel A indicates the average TM primary-market auction price and eBay secondary-market resale value for each concert-section-row in our matched dataset. Panel B indicates the TM face value and eBay 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 IV.

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.). The key difference is the payment rule: in TM's auction, winning bidders pay their own bid amount per ticket, independently of how high quality a ticket they won.[5] In contrast, in Google's "Generalized Second Price" position auction (GSP), winning bidders pay the next-highest bid amount times the number of clicks they receive, which in turn depends on how high a position they won.[6]

We provide a stylized theoretical analysis of TM's auction design, both to complement the empirical analysis and institutional detail and as a modest standalone contribution to auction theory. We adopt the modeling environment of Edelman, Ostrovsky, and Schwarz (2007) as well as their use of both a sealed-bid analysis and ascending-auction analysis; initially, the lone modeling difference is the payment rule. We begin by proving existence of an equilibrium in which the allocation is efficient: the solution concept is Bayes-Nash equilibrium for the sealed-bid model and perfect Bayesian equilibrium for the ascending auction model. These results then imply, via Myerson's Lemma, that the TM auction is revenue equivalent to

[5] TM described to us that the benefit of this payment rule is the simplicity of explaining it to customers.

[6] While Generalized *First* Price auctions have been heavily criticized in the context of keyword auctions (Edelman and Ostrovsky 2007), here the distinction between first-price and second-price per se seems smaller because the difference in quality between successive prizes is small—there are many pairs of tickets in the first row, many in the second row, etc. Rather, the larger source of strategic complexity in the TM auction is that bidders have to form a forecast of what quality of ticket they expect to win when determining their bid—their willingness to pay for a first-row ticket may be very different from that for a tenth-row ticket. We are grateful to Michael Schwarz for a helpful discussion about the relationship between TM's auction and other position auction formats.

GSP and to other efficient auction designs such as Vickrey-Clarke-Groves. They also immediately imply that the TM auction maximizes revenue subject to the constraint that all tickets are sold. Such a constraint could be interpreted in the context of Becker's (1991) observation that concerts that do not sell out may see a discontinuous drop-off in demand due to their social nature. We then add free entry by speculators to the analysis, in a stylized manner, and show that entry eliminates arbitrage profits for these bidders. We interpret this set of results—efficiency, constrained revenue maximization, and no arbitrage for resellers—as formalizing that TM's auction design is sensible.

That said, the equilibrium analysis suggests that bidding optimally in the TM auction is strategically complex. We find empirical 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. In the conclusion, we speculate that the complexity of participation in the auction may have been one factor that led to its demise.

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 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: both are nontrivial to price, both have histories of severe underpricing, both have histories of elaborate rent-seeking behavior associated with this underpricing (for IPOs, see, e.g., Nocera 2013), and both have secondary markets that are widely viewed to be efficient, suggesting that accurate pricing in the primary market is a realistic possibility.

Second, 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, Ostrovsky, and Schwarz's (2007) and Varian's (2007) studies of position auctions for keyword advertising.

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.[7] Paul Milgrom (Milgrom 2004; Section 1.4.1) has called this "one of the most frequent and misguided criticisms of modern auction design." Our study is a reminder of why 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

[7] This argument was made to one of the authors of the present study on his first day on the faculty at the University of Chicago, albeit over a friendly dim sum lunch.

ultimate allocation (see also Leslie and Sorensen 2014; Budish, Cramton, and Shim 2015; Budish, Lee, and Shim 2019).

The remainder of this paper is organized as follows. Section I provides institutional background. Section II describes TM's auction design. Section III presents the theoretical analysis. Section IV describes the data. Section V presents the main empirical results. Section VI compares experienced and inexperienced bidders and discusses potential modifications to TM's auction design. Section VII concludes by discussing why TM's auctions, despite "working" in the data, have failed to take off.

I. 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 recent prepandemic fiscal years, TM has sold about 500 million event tickets valued at on the order of $30 billion on behalf of clients including venues, promoters, sports leagues and teams, and museums and cultural institutions.[8] Tickets are typically sold at fixed prices that vary coarsely with seat quality; e.g., there might be just three or four 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. In 2011, TM estimated that 20 percent of all tickets purchased from TM in the primary market are subsequently resold on the secondary market (Ticketmaster Blog 2011a). Recent industry reports estimate that secondary-market dollar volume is on the order of $15 billion (Tan 2016; Duncan 2021). At the time of our data, eBay was the largest forum for secondary-market activity (Mulpuru 2008); at present the largest forum is StubHub ($4.7 billion in fiscal year 2019 annually (eBay 2019)), which has grown substantially since eBay's acquisition of it in 2007. TM 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 Entertainment LLC 2010). TM's secondary-market volume is presently on the order of $2 billion per year.[9] See Sweeting (2012) for a fascinating study of the dynamics of the secondary market, some findings from which manifest in our data as well; see further discussion in the online Appendix.

An important recent study by Leslie and Sorensen (2014) 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

[8] TM merged with Live Nation, a promoter, venue operator, and artist management firm, in January 2010. Live Nation's 2019 annual report indicates that TM sold 220 million tickets on which it earned per-ticket fees and another 267 million tickets through its clients' box offices for which it does not receive per-ticket fees (Live Nation Entertainment Inc. 2019). Earnings releases from 2016, 2017, and 2018 report similar numbers for total tickets sold and total gross transaction volume of between $28 and $33 billion.

[9] Live Nation's 2015 annual report indicated in its letter from the CEO that TM had $1.2 billion of secondary-market volume with growth of 34 percent year over year. Its 2016 and 2017 annual reports indicated growth rates of 24 percent and 16 percent, respectively, which imply secondary-market volume of $1.73 billion in 2017. Subsequent annual reports do not provide further information on secondary-market volume or growth.

is the rent-seeking activity that the possibility of resale encourages in the primary market. In Leslie and Sorensen's (2014) 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; Krueger 2019).

II. Ticketmaster's Primary-Market Auction

In 2003 Ticketmaster introduced auctions as a primary-market pricing method alongside fixed price. As discussed in the introduction, TM emphasized eliminating the arbitrage profits of Bob the Broker; the initial clients who adopted auctions also emphasized this idea that tickets are "worth what they are worth" and hence auctions are fair.[12] In this section we describe the TM auction in detail.

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's (2014) 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

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

[11] For instance, TM 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 Blog 2011b).

[12] Said Timothy J. Leiweke, former president of the Anschutz Entertainment Group, which produced the first event to use a TM auction: "Market inefficiencies…highlight the need for the event promoter to establish prices for live event ticketing closer to what the consumer is ultimately willing to pay. One way to establish a fair price for the best tickets is through an online auction, open to the general public, allowing the market to determine the price" (East Side Boxing 2003).

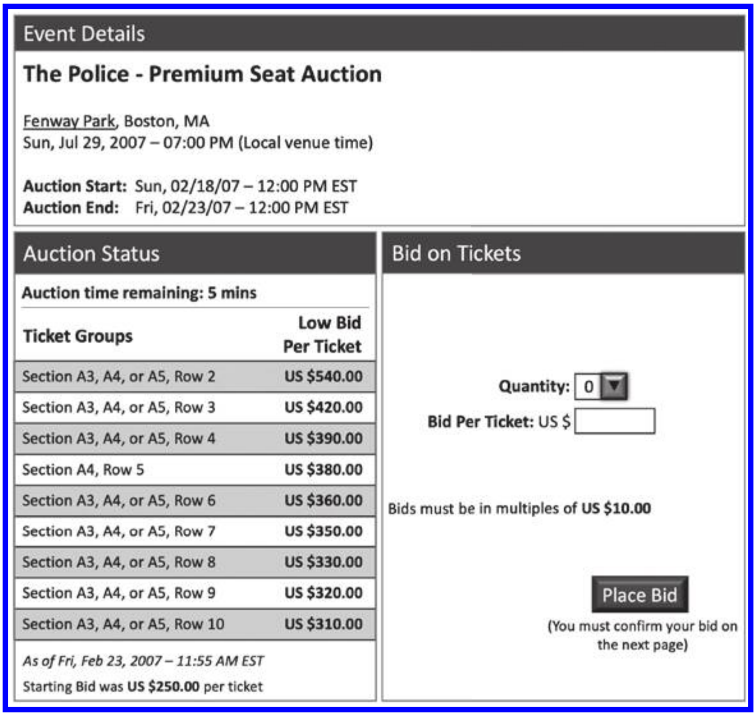

Figure 2. Illustration of the Ticketmaster Primary-Market Auction User Interface

Notes: This is a visual illustration of a real TM primary-market auction for tickets to see The Police in Fenway Park, Boston, July 29, 2007. For a description of the elements depicted, please see the text. The original screenshot can be seen in the working paper version of this paper, Budish and Bhave (2021).

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 preannounced 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 for this event went on sale on the intervening Tuesday.

The auction has a nonzero 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.[13] Bids consist of a per-ticket dollar amount and a desired number

[13] Auction bids are inclusive of all fees, whereas fixed-price tickets have separately stated face values and 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

of tickets—e.g., two or four. 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 three 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.

III. Theory Model of Ticketmaster's Position Auction

This section provides a stylized theoretical analysis of the Ticketmaster position auction, focused on the original 2003 rules. The purpose of the model is to clarify the relationship between the TM auction design and the position auctions used widely in internet advertising markets; to formalize that the TM auction design is "sensible" in that it has equilibria that satisfy attractive efficiency, revenue, and no-arbitrage properties; and to show that bidding in the auction is strategically complex and give a sense of what kinds of bidding mistakes may manifest in the data. We caveat that the model has important limitations. First, the model's notion of efficiency is purely allocative; that is, it does not account for an artist's distributional preferences (in contrast to Che, Gale, and Kim 2013 and Courty 2019). Second, the model's treatment of speculation is very stylized, as opposed to a structural model one could bring directly to the data as in Leslie and Sorensen (2014). Third, while the model does do a useful job of highlighting the strategic complexity of participating in TM's auction design, it is not rich enough to provide specific guidance on how to mitigate this complexity.

Readers primarily interested in the empirical analysis can skip directly to the next section.

Figure 2, the face value for tickets in the same quality tier as the auction was $225, and TM and venue fees were a combined $21.60, for an all-in fixed price of $246.60. 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 fees.

A. *Setup*

Our model closely mirrors the position auctions model of Edelman, Ostrovsky, and Schwarz (2007), in which there are vertically differentiated goods and goods are assigned in descending order of bid amounts. The key difference is the payment rule, which we clarify shortly.

There are K (pairs of) tickets and $n > K$ ex ante symmetric, risk-neutral bidders. 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.

Bidder i's private valuation for the k^{th}-best ticket is equal to $\alpha_k v_i$: v_i is bidder i's type, and α_k describes the quality of ticket k, with $\alpha_1 > \alpha_2 > \cdots > \alpha_K$.[14] 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, and normalize $\alpha_K = 1$.

We consider two ways to model the TM auction design, sealed-bid and ascending auction, analogously to how Edelman, Ostrovsky, and Schwarz (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 (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). Hence, we also consider an ascending auction model.[15] Our main results—on efficiency, constrained revenue maximization, and no arbitrage—obtain under both the sealed-bid and ascending auction approach.

In the *sealed-bid TM auction* model, each bidder submits a single bid. The bids are ranked in descending order, and then the k^{th}-highest bidder wins the k^{th} ticket, for each $k = 1, \ldots, K$. Winning bidders pay their bid amount; losers pay nothing. To explain the difference between this auction and 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. Here, the bidder's total payment is simply $b_{(k)}$: her own bid, without any adjustment for the realized quality.

In the *ascending TM auction* model, an auction clock is initialized at $p = 0$ and 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 $k = 2, \ldots, K$, gets the k^{th} ticket and pays the amount at which she herself dropped. The

[14] In Edelman, Ostrovsky, and Schwarz (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.

[15] A better model would combine an ascending auction phase and a sealed-bid phase, potentially also with learning about common value components of valuation along the way, but unfortunately this proved intractable.

$n - K$ bidders who do not get a ticket pay zero. This auction is related to the Generalized English Auction (GEA) of Edelman, Ostrovsky, and Schwarz (2007), analogously to how the sealed-bid TM auction is related to their treatment of GSP.

B. *Equilibrium, Efficiency, and Revenue*

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. Equilibrium is described as follows.

PROPOSITION 1 (Efficiency of Sealed-Bid Auction): *There exists a unique symmetric monotonic Bayes-Nash equilibrium of the sealed-bid TM auction in which all bidders bid according to*

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

The resulting allocation is efficient.

(All proofs are in the online Appendix.)

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.[16]

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

[16] While Proposition 1 is intuitively obvious ex post, and the bidding formula has relatively standard auction-theoretic intuition as described in the text, we note that it is not a foregone conclusion ex ante that novel auction formats have monotonic Bayes-Nash equilibria. For example, Gomes and Sweeney (2014) study Bayes-Nash equilibria of GSP auctions and find that they actually might not exist. (Edelman, Ostrovsky, and Schwarz 2007 use a complete information solution concept.)

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*

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

The resulting allocation is efficient.

The equilibrium of the ascending auction can be understood as follows. Given a bidder with value v, as long as there are at least K other active bidders in the auction, she behaves as if she is competing in a standard $K + 1^{st}$-price auction for K units of ticket K. That is, she simply bids up to her value for the K^{th} ticket of $v\alpha_K \equiv v$; that is, 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, $\frac{\partial T(v;\underline{v},k)}{\partial v}$, with their marginal benefit from doing so, $(\alpha_k - \alpha_{k+1})vkh(v)$. In the online Appendix, we prove by induction that the above collection of individual auction equilibria constitutes an equilibrium of the full ascending TM auction game.

Revenue.—By Myerson's Lemma (cf. Milgrom 2004), since the equilibria described in Section IIIB 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 TM auctions are revenue equivalent to any other efficient auction design in which the lowest type gets zero surplus, such as GSP or Vickrey-Clarke-Groves.*

Myerson's Lemma also implies the following corollary regarding the revenue performance of the TM auction:

COROLLARY 2: *If $v - (1 - F(v))/f(v)$ is strictly increasing, the sealed-bid and ascending TM auctions each maximize revenue subject to the constraint that all tickets are always sold.*

The constraint that all tickets are always sold can be interpreted in the context of Becker's (1991) observation that concerts that do not sell out may see a discontinuous drop-off in demand due to their social nature. An optimally set reserve price would increase revenue in our model but would risk leaving some tickets unsold.

Together, Propositions 1 and 2 and Corollaries 1 and 2 formalize that TM's auction has equilibria with attractive efficiency and revenue performance. The characterization of equilibria in Propositions 1 and 2 also suggest that playing TM's auction optimally is strategically complex.

C. *No Arbitrage*

We now 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 remainder are fans. Fans' types are drawn independently and identically from the continuously differentiable distribution $F_{fan}(\cdot)$ with support $[0,\bar{v}]$. Fix $\epsilon > 0$ and $w \in (\epsilon,\bar{v})$. Each reseller's type is drawn independently and identically from the continuously differentiable distribution $F_{pro}(\cdot)$ with support $[w-\epsilon, w]$. The interpretation is that w is the expected price (per unit quality) of a ticket on the secondary market, and resellers also face a small idiosyncratic cost in the interval $[0,\epsilon]$.

In each auction, n bidders are randomly drawn from the above population of potential bidders. This is equivalent to taking n iid draws from the distribution $F(\cdot)$, where

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

$F(\cdot)$ has support on $[0,\bar{v}]$, and we assume that $F(\cdot)$ is continuously differentiable $\forall \beta \in [0,1]$.[17] 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 $v \in [w-\epsilon, w]$ and any finite n_{fan},* $\lim_{n_{pro}\to\infty} s(v; n_{pro}, n_{fan}) = 0$.

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

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

win a ticket in the auction. Though a simple result, it highlights an important difference between auctions and fixed-price selling mechanisms.

IV. Data

Our data come from two sources: primary-market auction data provided by Ticketmaster and secondary-market resale value data scraped from eBay. Sections IVA and IVB describe each dataset in turn. Section IVC describes how the datasets are matched. The data are posted publicly with documentation at Budish and Bhave (2023).

A. *Primary-Market Auction Data*

Our primary-market auction data cover all TM auctions for concert tours that started in 2007. There are 22 concert tours, 576 concerts, and 759 auctions.[18,19] 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 two or four); time of bid; section, row, and seat numbers assigned to the bid; and 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.

B. *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 TM, 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 about 350,000 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

[18] For the majority of concerts, there is just a single auction, but in some cases there are two 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.

[19] We drop all concerts in Canada. These concerts comprise less than 2 percent of our primary-market dataset.

tickets being auctioned—specifically, section, row, and number of tickets.[20] We also observe eBay selling format parameters, such as the auction's opening bid and/or Buy It Now price[21] and the listing's total selling price. We then reduce this total selling price by eBay's transaction fees, which were roughly 4 percent at the time of our data[22] and then divide by the number of tickets to obtain a per-ticket net-of-fees selling price.

An example of an eBay auction web page is depicted in the working paper version of this paper, Budish and Bhave (2021; Figure 3). That eBay listing was for a pair of tickets in Section A3, Row 3 to see the Police at Fenway Park on July 29, 2007, and had a final auction price of $999.99, or $499.995 per ticket. This is the same concert whose TM auction web page we illustrated above in Figure 2, where tickets in Section A3, Row 3 had winning bids per ticket of as low as $420.00 and as high as $540.00 (i.e., the low successful bid in the next-highest-quality group). The screenshot may be helpful to look at for readers interested in how we are able to extract the necessary information from eBay listings to match that data with the TM auction data.

Unlike the primary-market data, we do not observe seat numbers in the eBay data.[23] Thus, when comparing the primary and secondary markets, we conduct our analysis at the section-row level, as described in Section IVC.

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 web page based on the number of tickets awarded per winning bid (typically one, two, or four) 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 percent of our eBay listings.

[20] For 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 our data began and subsequently renamed Item Specifics.

[21] A Buy It Now price is a price that, 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 Hortaçsu (2003, 2004).

[22] At the time of our data, eBay's fee schedule for Final Value Fees was 5.25 percent of the first $25, 3.25 percent of the amount between $25 and $1,000, and 1.5 percent of any amount above $1,000. 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 percent. 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 percent of the $580 transaction value.

[23] Listing tickets at the section and row but not seat level was a common practice on all of the major secondary-market websites at the time of our data and remains a relatively common practice to this day. See further discussion of this issue in Section IVC.

C. *Matching Primary- and Secondary-Market Data*

In this section, we describe our procedure for matching the TM 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.[24,25] 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 nonstandardized. For instance, eBay sellers typically use the string "1" in the Row field to describe tickets that are in Row 1, but we also observe entries of "#1," "**1**," "1st," "1 !!!!," "First," "one," "1 WOW!," and dozens of others. We handle this issue as follows. First, we create a dictionary that translates all observed eBay 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 TM 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. Our main specification for the analyses in Section V 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. For robustness, we also consider a specification that instead aggregates the TM transactions at the c-s-r level and treats each eBay observation separately, 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. the online Appendix). 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.

[24] In 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.

[25] While seat data being of negligible importance is typical, there are a handful of venues where heterogeneity in seat quality within a row is of sufficient importance that TM 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.

Table 1—Summary Statistics for Matched Primary- and Secondary-Market Data

	Full Ticketmaster dataset	Matched dataset	Percent 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 dataset	Mean	SD	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 IVA–IVC.

V. Do Primary-Market Auctions Discover Secondary-Market Values?

No Arbitrage.—Figure 1 (in the introduction) presents a scatterplot of our matched dataset at the level of the concert-section-row (c-s-r). In panel A, the horizontal axis denotes the average price per ticket in the Ticketmaster 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 (net of fees, as described in Section IVB). Panel B is identical, except that the horizontal axis denotes the tickets' face values. In both panels, the vertical distance between a point and the 45-degree line represents the profits associated with resale for that c-s-r.

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 potential resale profits, defined as the secondary-market resale value minus the primary-market auction price, are small. Figure 3 presents a histogram of these potential resale profits.[26] The mean potential resale profit is \$6.07, or 2.2 percent of the mean primary-market auction price of \$274.35. The 95 percent confidence interval of this estimate, clustering errors at the concert level, is $[-\$7.57, \$18.59]$;[27] 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 the online Appendix suggest that, if anything, the \$6.07 estimate is too high and mean potential resale profits are slightly negative.[28]

[26] To be clear, since we match the two datasets at the level of the concert-section-row but not seat, we are not claiming to observe any particular auction buyer's actual realized resale profits. Some buyers of course use the tickets themselves, and others may resell on secondary-market venues other than eBay. Rather, the spirit of the exercise is that we use eBay secondary-market values as indicative of the *potential* resale profits for a buyer who buys in TM's primary-market auction.

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

[28] Over our four matching specifications, the 95 percent confidence intervals admit estimates of potential resale 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 percent confidence interval for potential resale profits becomes $[-\$16.10, +\$10.86]$ under the main matching specification. For full details, see the online Appendix.

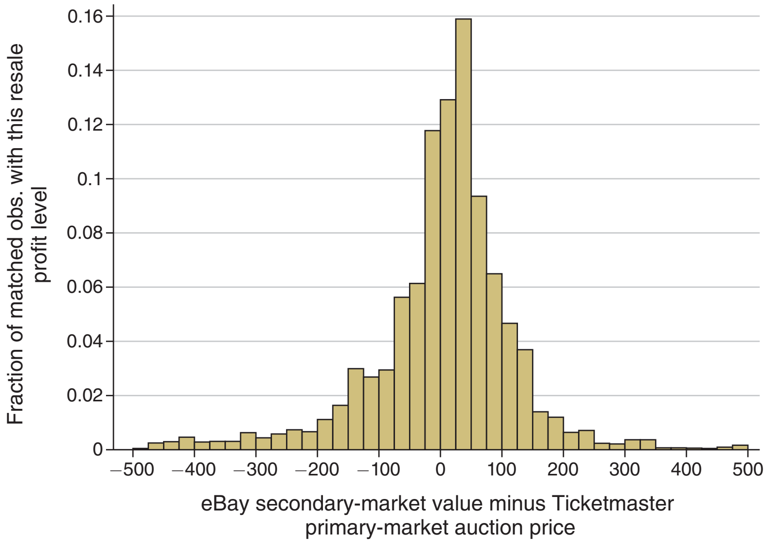

Figure 3. Distribution of Potential Resale Profits

Notes: Potential resale profits are calculated as the difference between the eBay secondary-market value and the TM primary-market auction price, averaged at the level of the concert-section-row. Profits are on a per-ticket basis and are net of eBay fees. For more details, see the text.

There are two other interesting features of the distribution of resale profits to highlight. First, there is substantial variance: while the mean potential resale 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, with a skewness of −0.69. 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 Section VIA and in the online Appendix.

Table 2, column 1 regresses the eBay secondary-market value on the TM primary-market price. The regression best fit line is not the 45-degree line in Figure 1, panel A but rather has 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 3—namely, the positive mode and the fat left tail.

Table 2, column 2 regresses the eBay secondary-market value on the TM primary-market face value. The interesting things to note are not the coefficients themselves but 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

Table 2—Price-Informativeness Regression Results

eBay secondary-market value	(1)	(2)	(3)
TM primary-market auction price	0.85		0.80
	(0.07)		(0.09)
TM face value		1.74	0.33
		(0.28)	(0.18)
Constant ($)	47.38	28.31	13.57
	(16.90)	(31.96)	(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.

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.[29] On average, the difference between the secondary-market resale value and the primary-market face value is $135.85, with a standard deviation of $215.24. In aggregate, over all of the tickets in our TM data, the TM auction raised $16.9 million of revenues, whereas the face values would have raised just $8.5 million.

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. Please see the online Appendix for robustness checks concerning the results in this section.

VI. Professional Resellers versus Ordinary Consumers

A. *Resale Profits*

Our results in Section VA 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, professional resellers may have specialized knowledge about which tickets to purchase, or be better at strategically bidding in the auction, than ordinary consumers. Hence, to fully assess whether the auctions eliminate the rents of Bob the Broker, we ideally would look separately at the arbitrage profits of professional

[29] Observe 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.

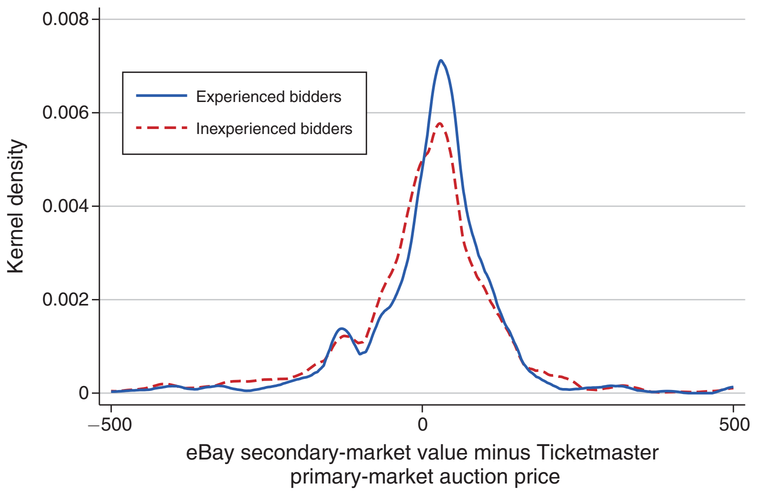

Figure 4. Potential Resale Profits, Experienced versus Inexperienced Ticketmaster Auction Participants

Notes: Experienced TM auction participants are defined as participants who win tickets in at least ten TM auctions. Potential resale profits are as defined in Figure 3. For more details, see the text.

resellers and ordinary individuals. 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.

While we cannot directly observe in our data whether a particular bidder is a professional reseller, we can exploit the fact that our TM data contain a unique bidder identifier to define a simple measure of experience in the auction—namely, the number of distinct auctions the bidder has won. We define a bidder as *experienced* if the bidder wins at least ten TM auctions (overall, not just restricted to the matched data). Such bidders account for 1 percent of the bidders in the TM data and roughly 16 percent of the transaction volume. We classify bidders who win between one and nine auctions as *inexperienced*. We think of auction experience as a proxy for being a professional reseller.[30]

Figure 4 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 percent level, with experienced bidders purchasing tickets with potential resale profits of

[30] Sixty-two percent 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). Twenty-three percent of volume corresponds to bidders with between two and nine transactions. We also consider a definition of experience based on the bidder winning at least two TM auctions in at least two cities (overall, not just restricted to the matched data). Such bidders account for 5 percent of bidders in the TM data and 24 percent of transaction volume. The results are very similar to our main specification. Last, we consider versions of both the 10+ auctions measure and the two-auctions-two-cities measure based on bids rather than wins. Again, the results move very little. See the online Appendix, sections B.A. and B.B.

$19.49 per ticket, while inexperienced bidders purchase tickets with potential resale profits of $2.47 per ticket.[31]

The figure also suggests that experience accounts for some of the asymmetry in the distribution of potential resale profits. Specifically, the most experienced bidders are significantly more likely to purchase tickets with small positive potential resale profits of between $0 and $100 per ticket (53.0 percent of transactions versus 42.4 percent of transactions, significant at 1 percent), and the most experienced bidders are significantly less likely to purchase tickets with losses that exceed −$100 per ticket—i.e., negative potential resale profits (11.7 percent of transactions versus 14.7 percent of transactions, significant at 5 percent). That is, the mode of small positive profits is disproportionately experienced bidders, whereas the left tail of large losses is disproportionately inexperienced bidders.

The fact that experienced bidders purchase tickets with 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.

B. *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. We find evidence of occasional severe overbidding for tickets in the highest-quality group in a given auction (e.g., first row): 14 percent of winning bids for tickets in the best-quality group are at least 25 percent higher than was necessary to win seats in that group, 5 percent are at least 50 percent higher than was necessary, and 1 percent of winning bids are at least 100 percent higher than was necessary.[32] Table 3 shows that these overbids are rarely submitted by experienced bidders and are disproportionately submitted by inexperienced bidders, especially for the overbids of at least

[31] We 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 these, section-row selection within a concert and paying a lower price for seats in the same section-row may reflect expertise in the auction per se (e.g., understanding that a bid that is provisionally winning for Row 1 may get bumped down to Row 2 if outbid), whereas artist and concert selection seem more likely to reflect superior information regarding what artist-city pairs will be highly demanded in the aftermarket rather than skill at the auction per se.

[32] 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 $1,000 would win tickets in the highest-quality group and pay substantially more than was necessary ($540 as of this screenshot), whereas any bid between $310 and $420 would pay an amount within $20 of the amount necessary to win tickets in the assigned row. In our data overall, 85.1 percent of winning bids are within $0–$10 of the next winning bid, and 94.1 percent are within $0–$50.

Table 3—Overbidding Analysis

Cutoff overbid percentage	25%	50%	100%	Proportion of full dataset
Proportion of overbids at least this large	0.14	0.05	0.01	($N = 6{,}125$)
Proportion of overbids this large that are				
Experienced bidders	0.06	0.02	0.00	0.15
	(0.01)	(0.01)	(0.00)	
Inexperienced bidders	0.94	0.98	1.00	0.85
	(0.01)	(0.01)	(0.00)	
First auction	0.75	0.78	0.76	0.69
	(0.02)	(0.03)	(0.06)	
Last auction	0.80	0.83	0.81	0.69
	(0.02)	(0.02)	(0.04)	
Only auction	0.65	0.68	0.63	0.56
	(0.02)	(0.03)	(0.06)	

Notes: Overbidding summary statistics for the winning bids in the highest-quality group in TM primary-market auctions. 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 percent (respectively, 50 percent, 100 percent). A bidder is defined as experienced if they win tickets in at least ten 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 titled "Proportion of full dataset" refers to all winning bids in the highest-quality group in the TM primary-market auction, not just overbids. Standard errors reported in parentheses are calculated using the bootstrap method. For more details, see the text.

50 percent (e.g., experienced bidders comprise 15 percent of the full dataset but only comprise 2 percent of the 50 percent overbids). 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.

It is not clear whether overbidding by inexperienced bidders should be viewed as a market design feature or a bug. 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 marketplace. Additionally, the risk of overbidding might deter some potential bidders from entering the 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; Azevedo and Budish 2019).

One way to mitigate strategic complexity would be to change the pay-as-bid element of TM's auction design to uniform pricing within row; e.g., all winners of tickets in Row k pay the highest winning bid for Row $k+1$. Call this the Generalized Uniform Price Auction. Another interesting idea is the auction design proposal of Sandeep Baliga and Jeff Ely (Baliga and Ely 2013a, b), called Purple Pricing, which features uniform pricing within each quality tier and is descending price rather than ascending price.

VII. Conclusion

This paper studies Ticketmaster's introduction in the mid-2000s of auctions into the primary market for event tickets. Our basic findings suggest that the auctions "worked": price discovery substantially improved; artist revenues roughly doubled versus the fixed-price counterfactual; and, perhaps most importantly, the auctions eliminated or at least substantially reduced potential resale profits for speculators. The only negative we found in the data was that inexperienced bidders made occasional large bidding mistakes, but this could be addressed by slight modification of the auction rules—e.g., to uniform pricing within each row or quality group (i.e., a Generalized Uniform Price Auction) or to the Purple Pricing auction design of Sandeep Baliga and Jeff Ely (Baliga and Ely 2013a, b).

And yet, over the decade that has passed since the time of the data, rather than coming into more widespread use, primary-market auctions for event tickets instead disappeared. LexisNexis searches suggest that TM auctions were in use from their introduction in 2003 through around 2011, with a peak in around 2005–2008 but that with limited exceptions, they have not been used since.[33]

We conclude by speculating as to why the auctions failed to take off. As discussed in the introduction, economic theory suggests that there are two basic choices for how to eliminate the rents of and rent-seeking by Bob the Broker: ban resale or set a market-clearing price. While auctions are no longer in use, what has at least partly taken off is using available data, including historical resale values, to set fixed prices in the primary market that more accurately approximate market clearing. An anecdote along these lines is the broadway show "Hamilton." We mentioned in the introduction that "Hamilton" adopted resale bans for 46 high-quality tickets per night, priced at just $10; they also sold an undisclosed number of high-quality tickets per night at $895 per ticket, a new broadway record, with that price chosen based on observed resale values at the time (Paulson 2016). More systematic evidence comes from examining pricing practices for Major League Baseball teams. Baseball teams host 81 home games per year and can use historical secondary-market data, historical sales patterns, and even knowledge about which opponents and pitchers are popular with fans to set prices for a particular game.[34] We hand collected data from baseball teams' websites and found that all 30 teams vary their prices by game at least somewhat, with on average 18 distinct pricing combinations (i.e., unique vectors of ticket prices) and on average a 2.2:1 price ratio between the most expensive and cheapest dates.[35]

[33] We did a LexisNexis search on the phrase "Ticketmaster" and then any one of a large number of phrases such as "ticket auction," "Ticketmaster auction," "premium seat auction," "premium ticket auction," "primary market auction," etc. and then hand checked articles and press releases for evidence of the use of TM's auctions. We found relevant articles and press releases in each year from 2003 through 2011 but zero relevant articles for the period 2012–2017. For the period 2012–2017 we also hand checked all articles with the words "Ticketmaster" and "auction" (a much wider screen) and found just three reported examples of TM auctions being used: the 2012–2013 Justin Bieber Believe Tour and two charity events. We also hand checked all summer 2017 TM concerts in arenas or stadiums in New York, Boston, Chicago, Los Angeles, and San Francisco and found zero concerts using auctions.

[34] The industry term for this practice is "dynamic pricing" (Goldstein 2012; Rishe 2012).

[35] We hand collected ticket prices from team ticketing websites during the last week of July 2017. Our data, which are meant to be illustrative, will understate the total amount of variation because they only contain about 40 percent of the baseball season. The average price ratio is computed by taking the unweighted average across all

We conjecture that the popularity of this practice relative to auctions partly reflects the simplicity and convenience for fans of posted prices relative to auctions, as has been documented more widely by Einav et al. (2018) and partly reflects a harder-to-model "repugnance" cost of ticket auctions (Roth 2007). It will be interesting to see if market designs along the lines of Purple Pricing, which have the flexible price discovery of an auction but aim to mitigate these harder-to-model negative aspects of auctions, ever come into more widespread use.

Some artists and events have indeed banned resale for their events, though this practice too remains relatively rare, and an entire 501(c)(4) lobbying organization, the Fan Freedom Project (initially funded by eBay and StubHub), is devoted to making the practice illegal (Lipka 2014; Budish 2019). 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 preteens, their parents, and several state attorneys general. For the artist's[36] next tour, Disney again set below-market prices but this time adopted technology that eliminated the possibility of secondary-market activity: just as airplane tickets are nontransferable because they are attached to the passenger's name, 2009 Miley Cyrus tickets were nontransferable because they were attached to a specific credit card that had to be presented in person at the concert venue (Waddell 2009). Other artists to have experimented with nontransferable tickets include Bruce Springsteen, Metallica, Justin Bieber, and U2 (Brooks 2017; Farhi 2010).

The 2012 Summer Olympics in London provided a cautionary tale regarding resale bans and may illustrate why they remain rare. 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 (unlike Hannah Montana fans). As a result, there were large blocks of empty seats at the Olympics, which was both wasteful and embarrassing for the event's organizers (Economist 2012). 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.[37] See Courty (2019) for a promising proposal in this direction, called "Centralized Exchange."

Setting market-clearing prices and banning resale are two ways to modify the primary market to eliminate Bob the Broker's rents. TM has also aggressively expanded into the secondary market, acquiring TicketsNow for $265 million in

types of tickets each team sells of the ratio of the maximum price to the minimum price for that type of ticket and then averaging across teams. The Chicago Cubs have the distinction of having a unique vector of ticket prices for every game observed in our data. The Chicago White Sox have the single most extreme example of price variation in our data, with a 12.6:1 ratio of the maximum to minimum price for the cheapest tickets in the stadium ($63 against the Chicago Cubs and $5 or $7 for all other games).

[36] A 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.

[37] Airlines accomplish this goal using a combination of refunds and cancellation fees (cf. Ely, Garrett, and Hinnosaar 2017); 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.

2008 (as well as UK-based Get Me In! for an undisclosed amount); entering into secondary-market partnerships with the National Basketball Association, National Hockey League, and National Football League (Major League Baseball has a partnership with StubHub); and most recently launching a secondary market within ticketmaster.com called Fan-to-Fan Resale that lists available primary-market tickets alongside secondary-market tickets.[38] This business exploits TM's unique ability, for events where it manages the primary market, to verify the authenticity of tickets in the secondary market. With transaction fees of about 30–40 percent in the largest secondary-market venues (Budish 2019)—of the full resale value, not of just the markup versus the fixed price—perhaps eliminating the rents of Bob the Broker is less profitable than taking a cut.

REFERENCES

Associated Press. 2007. "Ticketmaster Signs Resale Deal with N.F.L." *New York Times*, December 19. http://www.nytimes.com/2007/12/19/business/19ticket.html.

Azevedo, Eduardo M., and Eric Budish. 2019. "Strategy-Proofness in the Large." *Review of Economic Studies* 86 (1): 81–116.

Bajari, Patrick, and Ali Hortaçsu. 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics* 34 (2): 329–55.

Bajari, Patrick, and Ali Hortaçsu. 2004. "Economic Insights from Internet Auctions." *Journal of Economic Literature* 42 (2): 457–86.

Baliga, Sandeep. 2011. "How to Increase Revenue at Next Restaurant." http://cheaptalk.org/2011/04/11/how-to-increase-revenue-at-next-restaurant/ (accessed July 17, 2018).

Baliga, Sandeep, and Jeff Ely. 2013a. Big News. https://cheaptalk.org/2013/02/11/ big-news/ (accessed July 17, 2018).

Baliga, Sandeep, and Jeff Ely. 2013b. "The Power of Purple Pricing." *Harvard Business Review*, May 8. https://hbr.org/2013/05/any-business-trying-to-sell.

Becker, Gary S. 1991. "A Note on Restaurant Pricing and Other Examples of Social Influences on Price." *Journal of Political Economy* 99 (5): 1109–16.

Brooks, Dave. 2017. "After Rocky Start, U2 Sorts Out Paperless Ticket Issues on Joshua Tree Tour." *Billboard*, May 18. http://www.billboard.com/articles/business/7800921/u2-joshua-tree-tour-paperless-ticket-issues.

Budish, Eric. 2019. "Keynote Address, FTC Online Event Tickets Workshop." Keynote Address, FTC Online Event Tickets Workshop, Washington, DC, June 11, 2019.

Budish, Eric, and Aditya Bhave. 2021. "Primary-Market Auctions for Event Tickets: Eliminating the Rents of 'Bob the Broker'?" Unpublished.

Budish, Eric, and Aditya Bhave. 2023. "Replication data for: Primary-Market Auctions for Event Tickets: Eliminating the Rents of 'Bob the Broker'?" American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E161484V1.

Budish, Eric, and Lisa N. Takeyama. 2001. "Buy Prices in Online Auctions: Irrationality on the Internet?" *Economics Letters* 72 (3): 325–33.

Budish, Eric, Peter Cramton, and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics* 130 (4): 1547–621.

Budish, Eric, Robin S. Lee, and John J. Shim. 2019. "A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?" NBER Working Paper 25855.

[38] See Smith (2008a, b) for press coverage of the TicketsNow and Get Me In! acquisitions. See Associated Press (2007) and Jessop (2012) for press coverage of the secondary-market partnerships with the NFL and NBA, as well as http://www.ticketmaster.com/about/our-history.html for a corporate timeline of TM that indicates 2007 as the date of the three sports league partnerships. See Newman (2007) and ESPN (2012) on MLB's partnership with StubHub. See http://www.ticketmaster.com/verified for an overview of TM Fan-to-Fan Resale tickets. As discussed in Section I, TM's secondary-market volume is on the order of $2 billion annually.

Bulow, Jeremy, and Paul Klemperer. 1999. "The Generalized War of Attrition." *American Economic Review* 89 (1): 175–89.

Bulow, Jeremy, and Paul Klemperer. 2012. "Regulated Prices, Rent-Seeking, and Consumer Surplus." *Journal of Political Economy* 120 (1): 160–86.

Che, Yeon-Koo, Ian Gale, and Jinwoo Kim. 2013. "Assigning Resources to Budget-Constrained Agents." *Review of Economic Studies* 80 (1): 73–107.

Connolly, Marie, and Alan B. Krueger. 2006. "Rockonomics: The Economics of Popular Music." In *Handbook on the Economics of Art and Culture*, Vol. 1, edited by Victor A. Ginsburg and David Throsby, 667–719. Amsterdam: Elsevier.

Courty, Pascal. 2000. "An Economic Guide to Ticket Pricing in the Entertainment Industry." *Louvain Economic Review* 66 (2): 167–92.

Courty, Pascal. 2003a. "Some Economics of Ticket Resale." *Journal of Economic Perspectives* 17 (2): 85–97.

Courty, Pascal. 2003b. "Ticket Pricing under Demand Uncertainty." *Journal of Law and Economics* 46 (2): 627–52.

Courty, Pascal. 2019. "Ticket Resale, Bots, and the Fair Price Ticketing Curse." *Journal of Cultural Economics* 43 (3): 345–63.

Duncan, Byard. 2021. "How Is This Legal?" *Reveal*, March 8. https://revealnews.org/article/how-is-this-legal/.

East Side Boxing. 2003. "Lewis versus Johnson: Ticket Buyers Can Set Their Own Price for Best Seats Before They Go On-Sale." https://www.boxing247.com/news/st2504.php (accessed July 17, 2018).

eBay. 2019. *2019 Annual Report / Form 10-K*. San Jose, CA: eBay.

Economist. 2012. "Bring on the Touts." *Economist*, August 4. https://www.economist.com/leaders/2012/08/04/bring-on-the-touts.

Edelman, Benjamin, and Michael Ostrovsky. 2007. "Strategic Bidder Behavior in Sponsored Search Auctions." *Decision Support Systems* 43 (1): 192–98.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review* 97 (1): 242–59.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. 2018. "Auctions versus Posted Prices in Online Markets." *Journal of Political Economy* 126 (1): 178–215.

Ely, Jeffrey C., Daniel F. Garrett, and Toomas Hinnosaar. 2017. "Overbooking." *Journal of the European Economic Association* 15 (6): 1258–1301.

ESPN. 2012. "MLB Renews Deal with StubHub." *ESPN*, December 10. http://www.espn.com/mlb/story/_/id/8732189/major-league-baseball-stubhub-renew-secondary-ticket-market-deal.

Farhi, Paul. 2010. "'Paperless Ticketing' Aims to Thwart Scalping at Concerts, Sports Events." *Washington Post*, July 5. http://www.washingtonpost.com/wp-dyn/content/article/2010/07/04/AR2010070404180.html.

Friedman, Milton. 1991. "How to Sell Government Securities." *Wall Street Journal*, A8.

Goldstein, Jacob. 2012. "Baseball Teams Are Acting Like Airlines." *NPR*, March 6. http://www.npr.org/blogs/money/2012/03/05/147982995/at-the-ballpark-ticket-prices-change-as-fast-as-the-weather?.

Gomes, Renato, and Kane Sweeney. 2014. "Bayes-Nash Equilibria of the Generalized Second-Price Auction." *Games and Economic Behavior* 86: 421–37.

Jessop, Alicia. 2012. "The NBA and Ticketmaster Announce a Partnership Creating a New Ticketing Method." *Forbes*, August 20. https://www.forbes.com/sites/aliciajessop/2012/08/20/the-nba-and-ticketmaster-announce-a-partnership-creating-a-new-ticketing-method/#4a5d75254f7a.

Krueger, Alan B. 2019. "Rockonomics: A Backstage Tour of What the Music Industry Can Teach Us about Economics and Life." New York: *Currency*.

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

Leslie, Phillip, and Alan Sorensen. 2014. "Resale and Rent-Seeking: An Application to Ticket Markets." *Review of Economic Studies* 81 (1): 266–300.

Levitt, Steven D. 2008. "What Do I Have in Common with Hannah Montana?" www.freakonomics.com/2008/01/08/what-do-i-have-in-common-with-hannah-montana/ (accessed July 17, 2018).

Lipka, Mitch. 2014. "The Campaign against Paperless Tickets." *CBS News*, October 1. https://www.cbsnews.com/news/consumer-groups-campaign-against-paperless-tickets/.

Live Nation Entertainment, Inc. 2019. "Form 10-K." Beverly Hills, CA: Live Nation Entertainment, Inc.

Mankiw, N. Gregory. 2016. "I Paid $2,500 for a 'Hamilton' Ticket. I'm Happy About It." *New York Times*, October 21. https://www.nytimes.com/2016/10/23/upshot/i-paid-2500-for-a-hamilton-ticket-im-happy-about-it.html.

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

Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–1122.

Miranda, Lin-Manuel. 2016. "Stop the Bots from Killing Broadway." *New York Times*, June 7. https://www.nytimes.com/2016/06/07/opinion/stop-the-bots-from-killing-broadway.html.

Mortimer, Julie Hollard, Chris Nosko, and Alan Sorensen. 2012. "Supply Responses to Digital Distribution: Recorded Music and Live Performances." *Information Economics and Policy* 24 (1): 3–14.

Mulpuru, Sucharita. 2008. *The Future of Online Secondary Ticketing*. Cambridge, MA: Forrester.

Nelson, Chris. 2003. "Ticketmaster Auction Will Let Highest Bidder Set Concert Prices." *New York Times*, September 1. https://www.nytimes.com/2003/09/01/business/media-ticketmaster-auction-will-let-highest-bidder-set-concert-prices.html.

New York State Office of the Attorney General. 1999. *Why Can't I Get Tickets?: Report on Ticket Distribution Practices*. https://web.archive.org/web/20000816070216/http://www.oag.state.ny.us/press/reports/scalping/full_text.html (accessed May 26, 2011).

New York Times. 1867a. "Mr. Dickens in Boston – The Excess Demand for Tickets." *New York Times*, November 23. https://timesmachine.nytimes.com/timesmachine/1867/11/23/83466279.html?pageNumber=5.

New York Times. 1867b. "Mr. Dickens' Readings – Sale of Tickets for the Second Course." *New York Times*, December 12. https://www.nytimes.com/1867/12/12/archives/mr-dickens-readingssale-of-tickets-for-the-second-course.html.

New York Times. 1876. "A Carnival of Crime." *New York Times*, November 20. https://www.nytimes.com/1876/11/20/archives/a-carnival-of-crime.html.

Newman, Mark. 2007. "MLB.com Partners with StubHub." *MLB.com*. https://web.archive.org/web/20160222120127/http://m.mlb.com/news/article/2125070/.

Nocera, Joe. 2013. "Rigging the I.P.O. Game." *New York Times*, March 9. http://www.nytimes.com/2013/03/10/opinion/sunday/nocera-rigging-the-ipo-game.html.

Paulson, Michael. 2016. "'Hamilton' Raises Ticket Prices: The Best Seats Will Now Cost $849." *New York Times*, June 8. https://www.nytimes.com/2016/06/09/theater/hamilton-raises-ticket-prices-the-best-seats-will-now-cost-849.html.

Rishe, Patrick. 2012. "Dynamic Pricing: The Future of Ticket Pricing in Sports." *Forbes*, January 6. http://www.forbes.com/sites/prishe/2012/01/06/dynamic-pricing-the-future-of-ticket-pricing-in-sports/.

Rosen, Ellen. 2007. "In the Race to Buy Concert Tickets, Fans Keep Losing." *New York Times*, October 6. http://www.nytimes.com/2007/10/06/business/06money.html.

Rosen, Sherwin, and Andrew M. Rosenfield. 1997. "Ticket Pricing." *Journal of Law and Economics* 40 (2): 351–76.

Roth, Alvin E. 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives* 21 (3): 37–58.

Roth, Alvin E. 2008. "What Have We Learned from Market Design?" *Economic Journal* 118 (527): 285–310.

Roth, Alvin E., and Axel Ockenfels. 2002. "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.

Segrave, Kerry. 2007. *Ticket Scalping: An American History, 1850–2005*. Jefferson, NC: McFarland and Co.

Smith, Ethan. 2008a. "Ticketmaster Acquires U.K.-Based Get Me In." *Wall Street Journal*, January 28. https://www.wsj.com/articles/SB120157608344624355.

Smith, Ethan. 2008b. "Ticketmaster Buys Major Reseller." *Wall Street Journal*, January 15. https://www.wsj.com/articles/SB120036522352890281.

Smith, Ethan, and Sara Silver. 2006. "To Protect Its Box-Office Turf, Ticketmaster Plays Rivals' Tune." *Wall Street Journal*, September 12. http://online.wsj.com/article/SB115802470458860177.html.

Sweeting, Andrew. 2012. "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets." *Journal of Political Economy* 120 (6): 1133–72.

Tan, Gillian. 2016. "StubHub is Ebay's Next Ticket to Ride." *Bloomberg*, June 22. https://www.bloomberg.com/gadfly/articles/2016-06-22/stubhub-ebay-s-next-ticket-to-ride.

Ticketmaster Blog. 2011a. "An Open Letter to Fans and Partners on Ticket Reselling." http://web.archive.org/web/20111012043620/http://blog.ticketmaster.com/2011/08/12/an-open-letter-to-fans-and-partners-on-ticket-reselling (accessed July 4, 2018).

Ticketmaster Blog. 2011b. "The Ongoing Battle to Make Access to Tickets Fair for You." http://web.archive.org/web/20111012042211/http://blog.ticketmaster.com/2011/06/09/the-ongoing-battle-to-make-access-to-tickets-fair-for-you/ (accessed July 4, 2018).

Ticketmaster Entertainment LLC. 2010. "Fiscal Year 2009 Form 10-K." Beverly Hills, CA: Ticketmaster Entertainment LLC.

Varian, Hal R. 2007. "Position Auctions." *International Journal of Industrial Organization* 25 (6): 1163–78.

Waddell, Ray. 2009. "Miley Cyrus Fights Scalpers with Paperless Tickets." *Reuters*, June 21. https://www.reuters.com/article/us-cyrus/miley-cyrus-fights-scalpers-with-paperless-tickets-idUSTRE55K0CX20090621.

Zetter, Kim. 2010. "Wiseguys Plead Guilty in Ticketmaster Captcha Case." *Wired*, November 19. http://www.wired.com/threatlevel/2010/11/wiseguys-plead-guilty/.