Is Resale Needed in Markets with Refunds? Evidence from College Football Ticket Sales

Drew Vollmer*

February 24, 2024

Abstract

In what circumstances is allowing resale more efficient than providing refunds? I study common aftermarket policies in perishable goods markets with demand uncertainty. Using primary and resale market data on college football ticket sales, I estimate a structural model comparing resale, which has flexible prices but incurs frictions, to a partial refund scheme, which is centralized but has rigid prices. In the model, consumers anticipate shocks when making initial purchases, then engage in resale after shocks are realized. Because of resale frictions, refunds are more efficient on average. However, flexible prices make resale more efficient after large aggregate shocks.

*U.S. Department of Justice. Contact: drew.vollmer@gmail.com
The views expressed are those of the author and do not necessarily represent those of the U.S. Department of Justice. I would like to thank the editor, Andrew Sweeting, and three anonymous referees for valuable comments. I am also grateful to Allan Collard-Wexler, James Roberts, Curtis Taylor, Bryan Bollinger, Jonathan Williams, Daniel Xu, Juan Carlos Suàrez Serrato, Peter Newberry, Matt Leisten, and seminar audiences at Duke, the 2020 Southern Economic Association virtual meetings, and the 2021 DC IO Day. I would like to thank the university that provided its sales records for the project. All errors are mine.

1 Introduction

Economists often argue that resale is valuable because it reallocates goods to consumers with high values. For instance, a consumer might buy a concert ticket, learn she cannot attend, and resell to someone who can. But other methods of reallocating—like refunds—could reach the same result. With refunds, the consumer could return the ticket to the box office, then the primary market seller could sell it to someone else. In fact, many primary sellers, like airlines and hotels, offer partial refunds instead of allowing resale. Is resale better than alternatives like refunds?

In this article, the relative performance of resale and refunds depends on two main forces: aggregate demand uncertainty and resale frictions. The advantage of refunds is that sales are conveniently centralized in the primary market, but the cost is that all sales are made at the primary market seller's prices. If primary market prices are rigid, they will be suboptimal after aggregate demand shocks. In contrast, the advantage of resale is that prices in the resale market are flexible and adjust after shocks, but the cost is that decentralized resale markets incur frictions such as the hassle of browsing several markets. In markets with aggregate demand shocks and resale frictions, the optimal aftermarket design is ambiguous.

I carry out an empirical study to quantify the performance of resale and refunds and determine which design is best. Specifically, I develop and estimate a structural model of the market for college football tickets, which features rigid primary market prices, substantial aggregate demand uncertainty, and resale frictions. I use the model to conduct a counterfactual experiment in which resale is prohibited and the primary market seller offers a partial refund. The model predicts that partial refunds raise total welfare by 0.7% compared to resale, leave consumer welfare unchanged, and perform worse than resale when there are large aggregate shocks.

In the model, consumers purchase tickets over two periods. In the first period, a profit-maximizing monopolist primary market seller only sells a package of tickets to all games (season tickets). In the second period, it only sells tickets to individual games (single-game tickets). As in the data, the primary market seller has rigid prices. It sets all prices at the start of the first period and does not adjust them afterwards. Consumers resell or request refunds for tickets to individual games in the second period.

When choosing whether to buy season tickets in the first period, forward-looking consumers consider the aftermarket policy and have rational expectations of future demand shocks. There are two demand shocks, both realized between the two periods. First, consumers may receive idiosyncratic (independently drawn) shocks for each game. Idiosyncratic shocks are like schedule conflicts; they cause some consumers to have low values for their tickets and motivate reallocation. Second, all consumers receive a purely aggregate shock, like news about the team's quality, that shifts the demand curve for tickets. Because of the aggregate shock, the primary market seller's prices are likely to be suboptimal in the second period.

In the second period, consumers decide whether to participate in aftermarkets or buy tickets in the primary market. With refunds, consumers who bought season tickets can return any number of their tickets to the primary seller for a partial refund. With resale, they can resell any number of their tickets in the resale market. Consumers without tickets decide whether to buy tickets in the primary market or, when available, the resale market. To buy in the resale market, consumers must incur frictions that are not present in the primary market. Consumer decisions—and hence market outcomes like resale prices—depend on realized shocks. The structural model is needed to predict the full distribution of outcomes for circumstances not observed in the data, including different aggregate shocks and alternative aftermarket designs like refunds.

The analysis relies on a novel combination of data sets. The centerpiece is one season of primary and secondary market ticket sales records for a single university, with the secondary market records provided by the largest resale market, StubHub (Satariano, 2015). The records are informative about demand, the resale market, and the interaction between the primary and resale markets. I supplement the sales records with data on annual resale prices for 76 college football teams from 2011–2019 from SeatGeek, another online resale market. The average prices are informative about annual shocks to aggregate demand.

I estimate the model using the method of simulated moments, matching the model's aggregate predictions for quantities sold and resale prices to the data. Some parameters are estimated outside of the model simulations. These may be directly identified by the data, such as resale market fees, or estimated in reduced-form models, such as preferences for each game.

The main results come from counterfactual experiments where a profit-maximizing primary seller either allows resale or prohibits it and offers a partial refund. I find that resale performs better than refunds in states with large aggregate shocks. When the realized shock is one standard deviation above its mean, resale creates 10 percentage points more surplus in the second period relative to refunds than it does for an average shock. When the shock is one standard deviation below its mean, the corresponding advantage is five percentage points. Resale copes better because resale prices are flexible, rising \$9.49 above average when the shock is one standard deviation above average and falling \$8.55 when it is one standard deviation below. The difference in performance relies on the fact that model outcomes vary with realized shocks.

However, refunds perform better for less extreme realizations, and integrating over the distribution of demand shocks establishes that refunds are more efficient on average. I find that refunds raise total welfare by 0.7% without changing consumer welfare. Moreover, the primary seller prefers refunds because they raise profit by 2.4%. The changes imply that the harms of inflexible primary market pricing with refunds are outweighed by the elimination of resale market frictions. The welfare changes are meaningful relative to the number of tickets reallocated, 7.4% of all tickets sold with resale.

The analysis contributes to our understanding of aftermarkets and resale. There is little work comparing reallocation mechanisms. This article measures the performance of resale and refunds and ties it to resale frictions and aggregate demand shocks. In particular, the finding that resale frictions are large is important because it suggests integrated primary and resale markets—markets where primary and resale tickets are shown and sold together, which have become increasingly common—may enhance efficiency. The size of any benefits, however, depends on the precise reason for resale frictions, which this article cannot determine.

The comparison of resale to refunds also contributes to policy debates surrounding resale. The most prominent example is that governments have historically restricted or banned resale (Squire Patton Boggs LLP, 2017), but some now protect it (Va. Code §59.1-38.2). This article suggests that, at least in settings without systematically underpriced tickets, resale does not need to be protected if alternative methods of reallocation, like refund policies, are available.

Sports leagues' resale policies have also proven contentious. The NFL was sued for limiting resale on platforms other than Ticketmaster, its "official ticket marketplace" (Egelko, 2015). One possible justification for a single, official resale market is that it may reduce resale frictions, which this article finds are significant. However, this article does not study the potential benefits of competition between resale markets, which official marketplaces may dampen.

The analysis also offers suggestive evidence on alternatives to resale when there are rent-seeking brokers—a historic focus of the resale literature, as in Bhave and Budish (2023), Leslie and Sorensen (2014), and Courty (2019). Although underpricing and brokers are not a primary concern in this article, the model includes underpricing and rationing in equilibrium after high demand shocks. The results suggest that resale is significantly more efficient than refunds in such cases because it reduces rationing. However, the prediction is merely suggestive because costs associated with brokers, including waiting costs and market power, are not modeled.

The model has several limitations. The data do not identify brokers and so there are no brokers in the model. However, brokers are unlikely to be prominent in the market because tickets are not systematically underpriced for the university studied. A second limitation of the data is that ticket usage and resale outside of the online resale market are not observed. As a result, the model assumes that no one resells tickets to friends or coworkers and that all tickets sold are used.

The model assumes that the resale market is static with a single clearing price. The streamlined resale market is needed to compute a fulfilled-expectations equilibrium where outcomes vary with a distribution of shocks, but it is unable to capture dynamic pricing in the observed resale market. The model also assumes that consumers have homoegenous preferences for games and seat qualities, ruling out price discrimination based on seat quality and the heterogeneity that allows multiproduct bundling. Finally, I am unable to prove that equilibrium exists and is unique, although simulations reliably converge to a single equilibrium.

Related Literature. This article contributes to several literatures, most directly those on resale and demand uncertainty. For the literature on resale of perishable goods, this article estimates how resale affects profit and welfare by modeling both primary and resale markets. Leslie and Sorensen (2014) use a similar model combining primary and resale markets to study whether resale increases welfare in the market for concert tickets, but they do not consider profit because tickets are systematically underpriced in their sample. Tickets in my setting are not underpriced and so I study both profit and welfare. Sweeting

(2012) studies dynamic pricing in the resale market for sports tickets. Lewis et al. (2019) investigate the effect of resale on demand for season tickets in professional baseball but do not model how resale of season tickets affects sales of other tickets. The net effects of resale on buyers and primary market sellers are a traditional focus of the theory literature on resale, including studies such as Courty (2003) and Cui et al. (2014).

More broadly, this article relates to other studies of how to run aftermarkets for perishable goods. Two recent articles, Cui et al. (2014) and Cachon and Feldman (2021), have compared resale and refunds in theory, but neither conducts an empirical study or considers the effect of aggregate shocks.

Aftermarket design and resale have also been studied in the context of durable goods. With durable goods, primary market sellers compete against past vintages of their products, as in Chen et al. (2013). The durable goods problem leads to alternative aftermarket designs, such as leasing, studied in Hendel and Lizzeri (2002), and buybacks, studied in Hodgson (2023).

The current analysis also relates to studies of demand uncertainty in which aggregate uncertainty affects firms' strategic choices, such as Kalouptsidi (2014), Jeon (2022), and Collard-Wexler (2013). This article differs by focusing on strategies firms can use to cope with uncertainty.

The article relates to a significant literature on how to sell perishable goods. The choice of sales mechanism is studied in Waisman (2021), who considers whether sellers should use auctions or posted prices when reselling sports tickets, and in Bhave and Budish (2023), who consider the differences between auctions and posted prices for the very best concert tickets.

The most significant literature on selling perishable goods concerns dynamic pricing. Recent empirical studies include Lazarev (2013) and Williams (2022), who study dynamic pricing in airlines when there is demand uncertainty. Sweeting (2012) uses data on resale of sports tickets to determine which classes of theoretical models are most appropriate.

The primary seller's need to commit to prices in this article relates to another key issue in the dynamic pricing literature: whether the seller can commit to a pricing mechanism. The choice leads to different equilibria in Board and Skrzypacz (2016), where the seller can commit, and Dilmé and Li (2019), where the seller cannot and resorts to "flash sales."

2 An Example

In this section, I present an example illustrating that either refunds or resale can be most efficient, and that the key forces determining efficiency are demand uncertainty and resale frictions. The example provides a basis for the empirical model developed in Section 5, which has a similar structure. Numbered assumptions also apply to the empirical model (though they may be generalized), and I note any assumptions that do not apply.

A primary market seller has K tickets to sell over two periods. As detailed in Assumption 1, the primary seller has rigid prices.

Assumption 1. *The primary market seller announces a menu of prices $\{p_1, p_2\}$ at the start of the first period and cannot adjust it afterwards. Primary market tickets are sold at price p_t in period t.*

Each consumer i has value v_i for a ticket in period two,

$$v_i = (\nu_i + V)(1 - H_i). \tag{1}$$

Values have three components: consumer-specific tastes for tickets ν_i, a random variable V that affects all consumers' values, and a Bernoulli random variable H_i specific to consumer i. The realizations of the random variables are learned between the two periods, so consumer i only knows the expected value in the first period. Consumer values in the empirical model share the same structure but have additional terms.

The random variables H_i and V capture two sources of demand uncertainty. The first is strictly idiosyncratic and is described by the independently drawn Bernoulli random variable H_i, which equals one with probability ψ. Idiosyncratic shocks give each consumer a chance of unexpectedly being unable to attend the game—for example, if they have a schedule conflict—and make their realized value zero. The shocks cause some consumers who purchase tickets in the first period to reallocate in the second period. The second source of demand uncertainty is an aggregate shock, described by the random variable V. Aggregate shocks cause all consumers to have higher or lower values for tickets, like if the team performs better or worse than expected. In the example, there are two possible realizations of V: a high value V_H and a low value V_L. In the empirical model, V has a continuous distribution.

To simplify the example, I focus on the second period.[1] Suppose that the primary seller sells $\bar{Q}_1$ tickets in the first period with both resale and refunds. Because of idiosyncratic shocks, $\psi\bar{Q}_1$ of the tickets sold in the first period need to be reallocated. Thus $\bar{Q}_2 \equiv K - \bar{Q}_1 + \psi\bar{Q}_1$ tickets are available in the second period with both aftermarket policies. With refunds, all tickets are sold in the primary market, but with resale, $\psi\bar{Q}_1$ would be available in the resale market.

Also suppose the primary seller sets its price p_2 so that consumers exactly demand the remaining $\bar{Q}_2$ tickets in the second period when the aggregate shock V is realized as V_H. As discussed later, the conclusions are similar if p_2 is optimal when V is realized as V_L—the only requirement is that primary market prices be suboptimal for some realization of demand.

Refunds. Suppose that the primary market seller prohibits resale and offers a refund r.

Assumption 2. *When the primary market seller offers a refund, it prohibits all other ticket transfers. It selects a refund r at the same time that it selects prices. Consumers who purchase tickets in the first period can return them in exchange for r at the start of the second period, and any refunded tickets are returned to the primary market seller's inventory.*

For simplicity, suppose that the refund r is low enough that the only consumers who request refunds are the $\psi\bar{Q}_1$ who received schedule conflicts.[2] All $\bar{Q}_2$ remaining tickets are available in the primary market at the pre-selected price p_2. I illustrate the second period in Figure 1.

The supply curve S for tickets is horizontal at p_2 up to the number of tickets in the primary seller's inventory, $\bar{Q}_2$, and vertical afterwards.

The realized demand curve is either $D(p; V_H)$ or $D(p; V_L)$. (I suppress the first argument in the rest of the section.) For demand $D(V_H)$, the price p_2 is optimal and allocates all tickets. But when demand is $D(V_L)$, only $Q_2' < \psi\bar{Q}_1$ tickets are sold. The price rigidity thus creates deadweight loss in the low-demand state. The deadweight loss region can be split into loss from tickets sold in period one that were recovered through refunds (DWL_1) and loss on tickets unsold in period one (DWL_2).

[1] The conclusions in this section hold in a two-period equilibrium with forward-looking consumers. The empirical model features such an equilibrium.

[2] The assumption rules out the possibility that $0 < v_i < r$ for some consumers with tickets. This is possible if the realized V is low, and is permitted in the empirical model.

Resale markets. When there is a resale market, the primary market seller does not offer a refund. In the second period, consumers who bought in the first period can resell their tickets at prices of their choice; other consumers decide whether to purchase tickets in the primary and resale markets. I assume that the resale market is competitive, leading to an equilibrium with a single clearing price.

Assumption 3. *Participants take the resale price as given. An auctioneer announces a single resale price p_2^r that clears the resale market.*

Because of Assumption 3, the equilibrium resale price depends on realized demand and can be written as $p_2^r(V)$. The market clearing assumption ostensibly makes resale efficient, but resale-specific frictions temper any advantages. Specifically, consumers incur a friction of size s when purchasing in the resale market, reflecting factors like browsing costs and distaste for resale. Because of the friction, a consumer with value v_i is only willing to pay $v_i - s$ in the resale market. All consumers have the same friction s in this section, but frictions are heterogeneous in the empirical model.

The second period with resale is illustrated in Figure 2. Because consumers who bought early and received schedule conflicts are willing to accept any positive price, the supply curve is horizontal at zero up to $\psi\bar{Q}_1$. The supply curve is then horizontal at p_2 from $\psi\bar{Q}_1$ to $\bar{Q}_2$, reflecting the primary seller's remaining inventory. The supply curve then slopes upward because consumers with tickets and no schedule conflicts may be willing to resell at high prices.[3]

Because of resale frictions, consumers are willing to pay s less per ticket in the resale market at all realizations of V. The shaded region of height s at the bottom of the graph represents surplus lost to frictions in the resale market at both demand states. Owing to the effect of frictions on values, the prices at which supply and demand intersect overstate the resale price by s, as the axis labels indicate.

When demand is $D(V_H)$, all tickets are sold. When demand is $D(V_L)$, $\bar{Q}_2 - \psi\bar{Q}_1$ tickets are not sold, but all $\psi\bar{Q}_1$ tickets available for resale are sold at price $p_2^r(V_L)$.

Comparison. Because of resale's flexible prices, more tickets are sold in the low-demand state $D(V_L)$, raising welfare by up to DWL_1. But the gains are

[3] The final segment of the supply curve would be lower when demand is $D(V_L)$, but is not shown for simplicity.

tempered by losses to resale frictions incurred in both demand states. The net change in total welfare with resale is $\Pr(V = V_L) DWL_1$ minus *Frictions*. (DWL_2 is lost with both strategies when V_L is realized because the primary market price is p_2 in both cases.)

The net effect thus depends on the degree of aggregate demand uncertainty and the magnitude of resale frictions. Greater aggregate uncertainty, reflected in the distribution of V, leads to larger expected losses with refunds from mispricing in the primary market. In contrast, a larger friction s erodes the welfare gains from resale. The goal of the empirical exercise is to compare the two regions—DWL_1 and *Frictions*—in a model that features aggregate demand uncertainty, flexible resale prices, and resale frictions.

The ambiguity also applies to the primary market seller's profit. Price rigidities reduce profit with refunds, and frictions reduce the amount forward-looking consumers will pay for tickets with resale.

The conclusions of the example would be the same if realized demand had been unexpectedly high. If demand were higher than expected, there would be excess demand and inefficient rationing with refunds. Resale would lessen rationing, but would still incur frictions.

The empirical model includes an additional feature omitted from the example: resale fees. Resale markets are run by third parties who charge a fee on each transaction. The fee reduces the seller's profit with resale. As a result, it affects the profit-maximizing allocation in the first period—an effect not considered in this section because of the simplified first period—and the division of surplus in the second period.

3 Data

The analysis relies on a novel combination of two data sets. The first consists of ticket sales for a single university, covering both the primary and resale markets. Ticket sales are informative about demand for tickets and the extent of resale. The second consists of annual resale prices for football tickets at many universities, which are informative about year-to-year demand changes that reflect aggregate shocks to demand for each team.

Ticket Sales. The main data set includes primary and secondary market ticket sales for a large U.S. university's football team. The primary market records include ticket sales for two seasons. Each record indicates the price paid, date of purchase, and seating zone. Seating zones are clusters of seats sharing one price, which I use as a measurement of seat quality. The primary market records also indicate whether the sale was part of a season ticket package or promotion.

StubHub provides resale transaction records for the same university. The data do not include all resale because consumers also resell on competing sites. However, StubHub is likely to account for most resale because it was the largest resale platform at the time (Satariano (2015)). The university also requests that consumers resell on StubHub as part of a sponsorship deal.

The resale data do not include transaction prices. (The StubHub data was obtained through the university, and StubHub did not provide transaction-level prices to the university.) To learn about transaction prices, I use daily records of StubHub listings for the university's football games, gathered using a web scraper. The listing data only overlaps with the resale transaction data for the season studied in this article. Each listing includes a listing ID, price, number of tickets for sale, and location in the stadium (section and row).

The listing data scraped from StubHub do not directly reveal transactions, so I infer transactions from changes in listings. For example, if the number of tickets offered in a listing falls by two from one day to the next, then I assume two tickets were purchased at the last observed price.

The procedure leads to false positives because some listings are removed but not sold. I do not consider inferred transactions from the day of the game because many listings might be removed without a transaction. I correct for other false positives in two ways. First, I remove implausibly expensive transactions.[4] Second, the StubHub transaction data shows the true number of resale transactions at the game-section-time level, and I assume that the lowest-price inferred transactions are the true ones. The removed transactions are generally outliers, either occurring earlier or containing more seats than typical transactions.

Annual Resale Prices. I gather average annual resale prices for 76 college football teams from SeatGeek, another online resale market. The annual prices end in 2019 and start as early as 2011, although start dates vary by team. The SeatGeek data are informative about aggregate shocks. They show that the average price of a resold ticket varies meaningfully from one year to the next, reflecting changes due to shared factors like team performance.

The combination of data sets is novel. Although several studies have combined primary and resale market ticket sales (e.g. Leslie and Sorensen (2014)

[4]Transactions with prices over 1.5 times the 75$^{\text{th}}$ percentile of prices for similar quality seats.

and Bhave and Budish (2023)), to my knowledge, this is the first to add data on historic price fluctuations for similar events. The data on annual price variation make it possible to study the effect of aggregate demand uncertainty.

4 Descriptive Evidence

In this section, I provide descriptive evidence about the market that informs the empirical model.

Market Background. The university is a monopolist seller of its tickets in the primary market. There are other universities within driving distance, but they are not close substitutes because of local allegiances. In the season used in the analysis, the university sells tickets to five home games.[5] There are about 30,000 tickets available to the public for each game. (Other tickets are reserved for groups like students, athletics boosters, and visiting team fans.)

Tickets are sold in two main phases. The first consists of season ticket sales and takes place months before the season—80% of season tickets are bought at least four months before the season starts. The second phase occurs close to the game and consists of single-game ticket sales and resale. Single-game tickets do not go on sale until the first game is about a month away. 77% of resale and full-price single-game transactions occur after the season starts, and 50% occur within two weeks of the game. The empirical model reflects the timing of the market, with a first period in which season tickets are available and a second period in which single-game tickets and resale tickets are available.

The stadium is divided into five seating zones, which I use to measure the quality of each seat. Lower zones (e.g. zone 1) contain better seats, like those close to the field and near the 50-yard line.

Tickets available to the public are sold in several forms: season ticket packages; mini plans, which are bundles of tickets to a subset of games; and single-game tickets, including those sold at promotional rates. Of these, season ticket packages are by far the most popular, accounting for 75% of tickets sold.

The empirical model considers only season ticket packages and single-game tickets sold at full price to the public. Nonpublic tickets are excluded because they are off limits to most buyers. Mini plans are excluded because sales are

[5] An additional home game was scheduled but canceled. It is excluded from the data and so is also excluded from the analysis. I assume that consumers would have made the same season ticket purchases if that game had not been scheduled, and I use prorated season ticket prices in estimation.

negligible—they account for under 0.2% of tickets sold. Single-game sales at promotional and group rates are excluded because they are not optimally priced and may only be available to targeted groups, like veterans.[6]

The menu of primary market prices is shown in Table 1. Primary market prices mainly vary by seat quality. Prices vary slightly across games, but never by more than $10. Prices are not set so low that consumers anticipate sellouts, as they do for popular concert tours. Primary market tickets sold out for only one game in the season studied. In later seasons, the team performed better and a majority of games sold out.

Season tickets cost $25–$35 less than buying separate tickets to each game. Season ticket holders who are members of the university athletic club (which is not required to buy season tickets) receive additional perks such as reserved parking, access to a pregame tailgate, and the chance to visit practice. Variation in values for the perks is one reason consumers might have heterogeneous preferences for season tickets. However, I do not model athletic club membership because of a lack of data and because other sports at the university are more popular, providing a reason to join unrelated to football.

Resale Markets. Resale is a notable feature of the market, with 5.98% of all tickets sold to consumers resold on StubHub.[7] The overall rate of resale is higher because tickets resold on other resale markets are not observed. The rate at which consumers resell tickets will be used to estimate the frequency of schedule conflicts, captured in equation (1) as the probability ψ of activating the idiosyncratic shock term H_i.

The data support the idea that resale prices are flexible, reflecting the fact that resellers can adjust list prices at any time. Figure 3 demonstrates that resale prices are flexible because they adjust to differ from face value. It shows the distribution of face values and the distribution of the average fee-inclusive resale price for each game-quality combination. The differences reflect changes in demand, and the variation across games suggests that some games are more valuable.

Additionally, evidence from outside the market confirms that news leads to

[6] Nearly 40% of promotional tickets in the season were given away for free, and 98% were sold for half-price or less. Group tickets are discounted by over 40% on average. Promotions are not used to cope with demand uncertainty because they are too steeply discounted and too targeted.

[7] The figure is meant to reflect resale by consumers and so should exclude any sales to brokers. The primary market seller sells some tickets directly to brokers. I conservatively assume that all such tickets are resold on StubHub and remove them from both the numerator and denominator.

demand shocks and sharp changes in resale prices. For example, resale prices for Inter Miami rose by a factor of nine after Lionel Messi joined the team (De Avila, 2023).

The last important feature of resale markets is that they include fees and frictions that are not present in the primary market. The evidence for frictions is that consumers often buy tickets in the primary market when comparable resale tickets are cheaper. For instance, the average resale ticket to the first game is over $16 cheaper than the average primary market ticket, yet over 1,250 single-game tickets are sold in the primary market. There are several possible explanations for the friction. Consumers might not like or trust the resale market, they might find searching for tickets onerous, or they might be unaware of resale tickets. Frictions are reflected in Section 2 through the term s and are included in the empirical model.

StubHub charges a percent fee amounting to roughly 22% of the amount paid by the buyer.[8,9] The average combined fee is $10.71 on each ticket resold, compared to an average resale price under $40.

Annual Price Changes. SeatGeek's data on average annual resale prices for many universities provide information about aggregate demand shocks. The SeatGeek data provide observations on average resale prices $p_2^r(V)$ that can be used to learn about the underlying distribution of shocks V. To do so, normalize university u's price in year y by its average across all seasons,

$$\text{NormPrice}_{uy} = \text{AvgResalePrice}_{uy} / \left(\frac{1}{Y_u} \sum_y \text{AvgResalePrice}_{uy} \right), \qquad (2)$$

where $\text{AvgResalePrice}_{uy}$ is observed in the SeatGeek data and Y_u denotes the number of years in the sample for university u. Figure 4 shows the distribution of normalized prices for all teams after adjusting for time trends with a regression on year dummies. Year-to-year variation for each university is significant: the distribution is approximately normal and has an estimated standard deviation of .25, implying that there is a roughly one-third chance that prices in any given season will be more than 25% away from the mean. Further, dispersion

[8] Resale prices in this article are fee-inclusive to reflect the amount paid by the buyer.

[9] StubHub's fee structure is not public (StubHub, 2021). I use its typical fees, reported to be 15% of the fee-exclusive price from buyers and 10% from resellers (Goldberg, 2019). For a fee-exclusive price p, fees of $.25p$ when the buyer pays $1.15p$ imply a fee of 22% of the fee-exclusive price.

is not driven by a few outliers. The standard deviation of normalized prices is greater than .2 for more than 70% of all universities.

The dramatic swings in resale prices likely reflect aggregate demand shocks, such as changes in team performance that affect the common component of values V. For instance, in Clemson's lowest-priced season they lost two of their first three games—as many as they lost in the entire previous season—and prices were 30% lower than usual. In their highest-priced season they won the national championship and prices were nearly 35% higher.

5 Model

To reflect the data and match the theoretical framework, the model should feature advance purchases, resale frictions, aggregate demand shocks, and a resale market with outcomes that depend on aggregate shocks.

Outline, Utility, and Shocks

Let i index consumers. There are $j = 1, \ldots, J$ games in the season and $q = 1, \ldots, Q$ seat qualities. There are two periods.

The Primary Market. Two assumptions describe the behavior of the monopolist primary market seller.

Assumption 1′. *At the start of the first period, the primary market seller announces a menu of prices $\{p_{Bq}\}_{q=1}^{Q}$ for season tickets and $\{p_{jq}\}_{q=1}^{Q}$ for single-game tickets to each game j. Prices cannot be adjusted afterwards. Season ticket prices are only available in the first period and single-game prices are only available in the second.*

Assumption 1′ has the same content as Assumption 1 in Section 2, but allows different games and seat qualities. Price rigidities and the different phases of primary market sales reflect the setting, as discussed in Section 4.

Assumption 4. *A monopolist primary market seller maximizes profit, has capacity K_q for each seat quality $q = 1, \ldots, Q$, and has no marginal costs for each ticket.*

Assumptions 1′ and 4 define the monopolist's problem and hence the supply side of the market. There are no cost parameters, so the model does not rationalize observed prices in estimation—estimation only involves the demand side of the model. However, the assumption that the primary market seller maximizes profit matters for counterfactuals, where the primary seller sets different prices for each aftermarket policy.

Consumer Preferences. There are N consumers who want at most one ticket to each game. All consumers arrive in the first period, when they decide whether to buy season tickets or wait until the next period. In the second period, consumers who bought season tickets decide whether to resell tickets or attend each game. Consumers without season tickets decide whether to purchase in the primary market, secondary market, or not at all.

Consumer i's utility for a ticket of quality q to game j is measured in dollars (relative to an outside option normalized to zero) and has a similar form to equation (1) in Section 2,

$$u_{ijq}(V, H_{ij}) = \max\{\ \alpha_j\,(\nu_i + V + \gamma_q)\,(1 - H_{ij})\,,\ 0\ \}. \tag{3}$$

As before, the random variable H_{ij} captures idiosyncratic shocks like schedule conflicts and follows a Bernoulli distribution with success probability ψ. Consumer i receives independent draws of H_{ij} for each game and has no value for the game when H_{ij} equals one. The floor at zero reflects free disposal.

Similarly, the random variable V still acts as a common component of values for all consumers and so allows aggregate shocks, like an injury to the team's star player. Unlike in Section 2, the distribution of V is continuous. There is a single realization of V for the season.

Consumer i's utility also depends on a consumer-specific taste parameter ν_i and two new terms, a scalar α_j specific to game j and a scalar γ_q specific to seat quality q. The parameters α_j and γ_q do not vary among consumers, implying homogeneous preferences for games and seat qualities.

The effect of the parameters is most easily seen by looking at the terms in equation (3). The term featuring ν_i can be thought of as consumer i's base utility for a game. It is higher for better seat qualities q and higher consumer preferences for football games ν_i. The base utility is then multiplied by the game-specific α_j to explain why some games are more desirable.

Consumer heterogeneity in equation (3) is primarily vertical, ranking consumers by their value of ν_i. The ordinal ranking of consumers is fixed until idiosyncratic shocks H_{ij} are realized, but the cardinal differences may be stretched across games (through α_j) or shifted by qualities and shocks (through

V and γ_q). Heterogeneity in game and quality preferences is not essential to the research question and would make computation more challenging.

Equation (3) does not have a traditional error term, so consumer i's expected values are perfectly correlated across games j. The model includes other features that create random variation in choices over the products in a consumer's choice set, such as season tickets and resale tickets.

The parameters ν_i and V require parametric distributions.

Assumption 5. *The common value V follows a normal distribution, $V \sim N(0, \sigma_V^2)$.*

Assumption 6. *The consumer taste parameter ν_i follows an exponential distribution, $\nu_i \sim Exp(\lambda_\nu)$.*

Assumption 5 follows from the distribution of normalized prices in Figure 4. Assumption 6 is needed to limit the number of parameters.

Model Outline. The model outline is depicted in Figure 5. Consumer decisions in period two are depicted for a single game j but occur for all games.

The rest of the section describes the choices detailed in Figure 5. Using backward induction, I start in period two.

Period Two

At the start of period two, consumers know whether they bought season tickets and learn the realizations of shocks: schedule conflicts H_{ij} and the common value V. Consumers with season tickets decide whether to resell or attend. All other consumers decide whether to purchase tickets in the primary market, resale market, or not at all.

The resale market operates as described in Assumption 3. Specifically, there is a schedule of resale prices $\{p^r_{jq}(V)\}_{q=1}^{Q}$ clearing the resale market for each game j, all resale market participants act as price takers, and the resale price varies with the realization of V and includes all fees paid by buyers. The assumptions simplify the search for resale prices, which must be conducted for each realization of V. The assumptions do not imply a perfectly competitive market that maximizes welfare gains: features like resale frictions reflect the data and reduce the efficiency of resale. The market-clearing price also implies that no resellers have market power, which may be incorrect if there are large brokers.

Consider the supply side of the resale market for game j. Consumers who bought season tickets resell if the proceeds exceed their utility from attending,

$$u_{ijq}(V, H_{ij}) \leq (1-\tau)p^r_{jq}(V). \tag{4}$$

The resale fee τ is a percentage of the resale price, matching the policies of resale markets like StubHub. The condition implies that consumers who receive an idiosyncratic shock are willing to resell at any positive price.

Consumers without season tickets decide whether and how to buy tickets to game j. They have three choices: make no purchase and receive surplus zero (noted as *No Purch. Surplus*$_{ij}$), purchase in the primary market and receive surplus *PM Surplus*$_{ijq}$, or purchase in the resale market and receive surplus *SM Surplus*$_{ijq}$. The surplus terms are

$$\textit{No Purch. Surplus}_{ij} = 0, \tag{5}$$

$$\textit{PM Surplus}_{ijq}(V, H_{ij}) = u_{ijq}(V, H_{ij}) - p_{jq}, \tag{6}$$

$$\textit{SM Surplus}_{ijq}(V, H_{ij}, s_{ij}) = u_{ijq}(V, H_{ij}) - p^r_{jq}(V) - s_{ij}, \tag{7}$$

where s_{ij} is a friction that affects the surplus from buying in the resale market.

The friction explains why consumers might purchase primary market tickets when similar tickets are cheaper in the secondary market. I make several assumptions about the friction for tractability.

Assumption 7. *The frictions s_{ij} follow an exponential distribtion, $s_{ij} \sim Exp(\lambda_s)$, and are independently drawn for each individual and game. Consumers know the distribution of frictions in the first period but do not learn their realizations until the second.*

For each game, each consumer without season tickets chooses the option maximizing surplus among equations (5), (6), and (7). Together, the equations determine demand for primary and resale market tickets in period two.

Consumers may not be able to choose their surplus-maximizing option if an alternative sells out, which is plausible in states with a high value of V. In such cases, I assume that tickets are rationed randomly.

Assumption 8. *Tickets are rationed randomly when there is a stock-out. The probability of receiving a primary market ticket of quality q to game j is $\sigma_{jq}(V)$.*

Period One

In period one, consumers decide whether to buy season tickets based on expected outcomes in period two. The model includes a decision rule reflecting the ability to resell in the second period, which is modified in counterfactuals with other aftermarket policies.

Consumers with season tickets ultimately receive the maximum of their value for attending game j and the after-fee resale price. Consumer i's surplus from season tickets of quality q is

$$\begin{aligned} ST\ Surplus_{iq} = \sum_j \mathrm{E}_{V,H_{ij}} \Big(\max \big\{ u_{ijq}(V, H_{ij}), \\ (1-\tau) p^r_{jq}(V) \big\} \Big) + \delta_i - p_{Bq}. \end{aligned} \tag{8}$$

Surplus depends on the price of season tickets and an additional parameter δ_i, which captures heterogeneity in consumer values for season tickets. Heterogeneity could come from the effects of attending many games or perks for season ticket holders in the university's athletics club. Consumers sharing the same ν_i could make different season ticket decisions because of variation in δ_i.

Assumption 9 limits the season ticket preferences δ_i to two values, providing a parsimonious way to describe consumer heterogeneity.

Assumption 9. *The parameter δ_i satisfies $\delta_i \in \{\delta_L, \delta_H\}$, where $\delta_L < \delta_H$. Values of δ_i are independently drawn. A fraction ζ of consumers have $\delta_i = \delta_H$, implying* $\Pr(\delta_i = \delta_H) = \zeta$.

The parameters δ_H and δ_L may not be point identified. If no consumers of type L want to buy season tickets at some value of δ_L, they also would not buy season tickets at any lower value of δ_L. As discussed in Section 6, δ_L is partially identified in the estimated equilibrium.

The surplus from waiting until period two requires an expectation for each game j's surplus. Consumer i's set of alternatives for game j is

$$\begin{aligned} \mathcal{C}_{ij}(V, H_{ij}, s_{ij}) = \{0, \{SM\ Surplus_{ijq}(V, H_{ij}, s_{ij})\}_{q=1}^{Q}, \\ \{PM\ Surplus_{ijq}(V, H_{ij})\}_{q=1}^{Q}\}. \end{aligned} \tag{9}$$

A consumer's surplus from waiting, $WaitSurplus_i$, is the expected value of equation (9) after accounting for rationing.[10] The consumer's choice set in

[10]Complete expressions can be found in the web appendix.

period one is thus

$$\mathcal{C}_{i,ST} = \left\{ WaitSurplus_i, \{ST\ Surplus_{iq}\}_{q=1}^{Q} \right\}. \tag{10}$$

Each consumer selects the maximizer of equation (10). If there is a stock out, tickets are again rationed randomly.

Equilibrium

I search for a fulfilled-expectations equilibrium. The primary market seller anticipates consumer demand and selects profit-maximizing prices $\{p_{Bq}\}$ and $\{p_{jq}\}$. (Equivalently, the primary market seller maximizes revenue because the marginal cost of a ticket is zero.) Consumers know the equilibrium resale price functions $\{p^r_{jq}(V)\}$ and primary market purchase probabilities $\{\sigma_{jq}(V)\}$. Consumers make optimal choices in the first period given expectations for resale prices and probabilities, and their expectations are realized in the second period when they make optimal purchase choices.

Although I am unable to prove that an equilibrium exists,[11] model simulations reliably converge. Simulations do not suggest there are multiple equilibria.

The search for equilibrium is computationally demanding, requiring repeated iteration to find the fixed point in resale price and primary market rationing functions. Iteration is time-consuming because there are $2J$ such functions, one of each type for each game. I discretize the distributions of values ν_i and common values V. The standard for convergence is that resale prices are within \$0.01 between iterations for any realization of V, and the mean primary market purchase probability for each quality is within one percentage point. It takes over 20 minutes to converge for the optimal demand-side parameters described in Section 6.

6 Estimation and Results

Estimation focuses on the demand side of the model because there are no supply-side parameters to estimate.

[11] The challenge in applying familiar fixed-point theorems is that the domain and codomain may not be compact. When few season tickets are sold, resale prices could be arbitrarily high. Further, the arguments to the function whose fixed point we seek are functions ($p^r_{jq}(V)$ and $\sigma_{jq}(V)$).

K_q, τ, and ψ. These parameters are identified directly from the data. I use the university's designations for seating zones q and take K_q to be the number of seats in zone q. The fee τ is the percentage of the fee-inclusive price paid by the buyer, calculated directly from StubHub's policies.

The probability ψ of receiving an idiosyncratic shock (and having H_{ij} equal one) is identified by the frequency of resale. In the model, observed resale is explained by idiosyncratic shocks in equilibrium, so the parameter ψ equals the ratio of tickets resold by consumers to all tickets sold. However, the true number of tickets resold is unknown because some consumers resell on other platforms, which are not observed. I conservatively assume that StubHub is 75% of the resale market. Leslie and Sorensen (2014) assume a 50% share for StubHub and eBay and Satariano (2015) reports that StubHub has roughly half of the ticket resale market.

α_j *and* γ_q. The parameters α_j and γ_q affect consumer values and hence resale prices. They are recovered from a reduced-form model of resale transaction prices. The price of resale listing k may vary because of characteristics X_k that are not otherwise modeled, such as the number of tickets in the transaction and the number of days until the game.

Assumption 10. *The resale price p^r_{jqk} of transaction k follows*

$$p^r_{jqk} = \alpha_j(\gamma_q + X_k\beta) + \varepsilon_{jqk}. \tag{11}$$

Equation (11) is similar to the expression for consumer values in equation (3), but also allows the characteristics of listing k to affect resale prices. The key assumption is that changes in consumer values pass through completely to resale prices. The counterfactual results in Figure 7 suggest that the assumption holds in equilibrium. For more discussion, see the web appendix.

I estimate the model using nonlinear least squares and obtain standard errors using the bootstrap.

The identifying variation for α_j and γ_q comes from across-game and across-quality variation in resale prices. More precisely, α_j explains why similar tickets for different games sell at different prices and γ_q explains why tickets to the same game with different qualities sell at different prices.

σ^2_V. The parameter σ^2_V is the estimated variance of a normal fit to the set of aggregate demand shocks $\{V\}$. The set of shocks is not directly observed and must be constructed from the data.

The basis for the set of shocks is observed variation in annual resale prices for the teams in the SeatGeek data. I limit the SeatGeek data to similar teams, defined as universities in similar athletic conferences that—like the university studied—have not appeared in a national championship game or the College Football Playoff in the past 20 years.

The starting point is the normalized resale prices described in equation (2). The normalized prices give university u's percent deviation in resale prices in year y from university u's average across all seasons. I use the normalized prices to obtain a set of all observed shocks $\{V_{uy}\}$.

Assumption 11. *The shock V_{uy} realized for university u in season y is*

$$V_{uy} = \bar{p}^{r}_{\bar{u}} \left(NormPrice_{uy} - \overline{NormPrice}_{y} \right) \left(\sum_{j} w_j \alpha_j \right)^{-1} . \tag{12}$$

I obtain σ^2_V by fitting a normal distribution to the set $\{V_{uy}\}$. Equation (12) multiplies the normalized price (a percent deviation) by the studied university $\bar{u}$'s mean season-adjusted resale price, $\bar{p}^{r}_{\bar{u}}$, to obtain an absolute price shock. The equation also adjusts for time trends, through $\overline{\text{NormPrice}}_{y}$, and the weighted average value of $\alpha_j V$. The weights w_j are taken as game j's share of all resale transactions. Because estimation is based on normalized prices, the identifying variation comes from season-to-season changes within each university.

The procedure makes three assumptions. First, the year-to-year variation in the SeatGeek data is the sole source of variation in the common value. The assumption could understate the variance because of game-specific shocks like rain, but it could also exaggerate the variance if the year-to-year change is predictable, like when a star player graduates. Second, shocks to the common value pass through completely to resale prices. Like for the estimation of α_j and γ_q, Figure 7 suggests the assumption holds in equilibrium. And third, the university faces the same shocks to normalized prices as other schools.

λ_s, λ_ν, δ_H, δ_L, and ζ. The remaining parameters govern the distribution of resale frictions (λ_s), the distribution of values (λ_ν), preferences for season tickets (δ_H and δ_L), and the fraction of consumers with each value of δ_i (ζ). They are estimated using the method of simulated moments (MSM).

For each set of candidate parameters, I take observed primary market prices as given and simulate the demand side of the model to predict 11 total moments:

the number of season tickets sold, the average resale price for each of five games, and the quantity of primary market tickets sold for each of five games. MSM with aggregate moments is necessary because there are no closed-form expressions for model predictions as a function of parameters, and because many individual-level choices, such as to not purchase, are not observed.

An important detail for estimation is that the observed moments in the data result from the single value of V realized in the season studied. I determine the value of V realized in the data using equation (12), which compares observed to average resale prices for the university. Then I estimate the model by comparing model predictions for the realized value of V to the observed data.

Average resale prices in the data depend on the composition of seat qualities resold. I weight model-predicted resale prices, which are at the game-quality level, with the observed quality composition for resold tickets in the data to ensure they are comparable.

The weight matrix has zeros off the diagonal and treats a 1% deviation in each moment from its observed value equally. (The inverse covariance matrix cannot be recovered because estimation moments come from separate datasets.)

In simulations, I discretize the distributions of ν_i and V. The grid of values for ν_i ranges from 0 to 198 in increments of 0.1. The grid for V consists of 100 values, the evenly spaced quantiles of the distribution from 0.5% to 99.5%. The extreme values of V in the grid are ± 20.8. There are $N = 200{,}000$ consumers.

I calculate standard errors for the parameters estimated with MSM using the bootstrap. I draw a sample of 50 estimation moments, then estimate optimal parameters for each set with the other parameters fixed at their point estimates. The reported standard errors are derived from the set of estimated optimal parameters and do not account for uncertainty over previously estimated parameters.

To obtain the set of alternative estimation moments, I calculate each moment's variance and sample from the implied distributions. For each game's average resale price, I construct the variance by sampling from the distribution of transaction-level resale prices. For the number of season tickets and single-game primary market tickets sold, I do not observe individual-level choices and so cannot sample them. Instead, I assume that the underlying individual choices are based on Bernoulli draws and use the associated variance. For details, see the web appendix.

Each parameter is identified by a combination of the estimation moments.

Start with the distribution of resale frictions, parameterized by λ_s. In the model, consumers purchase in the primary market if the primary market price is less than the sum of the resale price and the resale friction. For instance, if the resale price is \$5 less than the primary market price, any consumer with $s_{ij} > 5$ prefers the primary market. The distribution of s determines the number of consumers with $s_{ij} > 5$ and hence the number of tickets sold in the primary market. It follows that λ_s is identified by primary market quantities and the difference between primary and resale market prices.

Next, consider the distribution of values for college football relative to the outside option, parameterized by λ_ν. Higher values cause primary market quantities and resale prices to rise, so λ_ν is explained by all estimation moments: season ticket quantities, primary market quantities, and resale prices.

Finally, consider the parameters related to season tickets, δ_H, δ_L, and ζ. They relate to season ticket purchases. The parameters either directly raise values for season tickets (δ_H and δ_L) or increase the fraction of consumers with high values for season tickets (ζ). However, the quantity of season tickets sold is not enough to identify all three parameters. Raising δ_H while lowering δ_L or ζ could result in the same quantity of season tickets sold.

Resale prices and primary market quantities provide the additional information needed to identify the season ticket parameters. By determining which consumers buy season tickets, the season ticket parameters also determine which consumers remain in the second period to form the demand curves for tickets in the primary and resale markets. For example, suppose that δ_H were raised and ζ lowered from their true values so that the number of season ticket buyers stayed the same. The new high-type season ticket buyers would have lower values ν_i than before, so buyers in the second period would have higher values than before. The change would affect resale prices and primary market quantities.

As mentioned in Section 5, the parameters δ_H and δ_L are partially identified in general: a range of values produces the same allocation and model fit. Partial identification does not impede the estimation procedure. It is sufficient to find any parameters producing the best-fit allocation, then evaluate whether a higher δ_L or lower δ_H would produce the same allocation when other parameters are held fixed.

Results and Fit. Estimated parameters are in Tables 2 and 3. As noted in Section 4, the resale fee is about 22% of the fee-inclusive price paid by the buyer. The idiosyncratic shock rate suggests that 8% of buyers change their minds about attending the event between the first and second periods.

Consumer values vary widely across games and qualities. I normalize $\alpha_1 = 1$ and $\gamma_5 = 0$. The most valuable game has attendance values 67% higher than those for the baseline game; the least valuable game has values nearly 50% lower. The best seats are worth almost $23 per ticket more than the worst seats for game 1, with the difference scaled by the relevant α_j for other games.

The standard deviation of the distribution of consumer values is $8.08. The university thus faces consumer values for the baseline game that differ from the mean by more than $8 about a third of the time.

In Table 3, the parameter defining the exponential distribution of resale market frictions is 78.05. Hence the mean consumer has a friction of $78.05 associated with buying resale tickets, but there is substantial dispersion. Over a quarter of consumers have draws of $25 or less. Nonetheless, the friction is substantial and reduces the efficiency of the resale market.

The mean of the distribution of consumer types is 17.60, suggesting that the average consumer would pay $17.60 for the worst seats to the baseline game in an average season.

Consumers with high values for season tickets are estimated to represent 81% of the population and value season tickets $5.60 more than buying tickets to each game separately.

The corresponding estimate for consumers with low values for season tickets is partially identified. At the optimal parameters, the model predicts that no consumers with low values for season tickets will buy a season ticket package. Holding the other parameters constant, the model fit and allocation are the same for any δ_L in $(-\infty, -203.06]$. I report the upper bound of the set and base standard errors on the upper bounds of the identified sets for the alternative moments. The magnitude of δ_L ensures that some consumers with high values demand tickets in the second period.

Partial identification is not consequential because the results of counterfactual experiments do not vary within the identified set.

Table 4 and Figure 6 assess the model fit, where second-period model predictions are for the value of V observed in the data. Observed and model-implied resale prices are extremely close. The model captures the patterns in primary market sales across games but does not fit them exactly. The looser fit is expected because there are no parameters specifically designed to fit game-specific quantities. Finally, the model-implied number of season tickets purchased is within 10% of the true value.

Optimal prices from the resale counterfactual in Section 7 can be compared to observed prices, which the model does not attempt to rationalize. Optimal prices are about $9 higher than observed prices on average, and vary considerably more across games.

7 Counterfactuals

The estimated model makes it possible to evaluate several counterfactual policies. In addition to the main experiment on partial refunds, I implement counterfactuals to measure the effects of market features like primary market price rigidities and resale fees.

The model predicts allocations, welfare, and profit for each realization of the aggregate shock V. Reported counterfactual results are average outcomes obtained by integrating out over the distribution of shocks.

Unlike the estimation procedure, the counterfactuals use the assumption that the primary market seller maximizes its profit. In each counterfactual, I solve for the primary market seller's optimal menu of prices and evaluate welfare at those prices. Solving for optimal prices is necessary because the aftermarket policy affects the optimal price menu.

I place a mild restriction the primary market seller's choice of prices to simplify the search for profit-maximizing prices.

Assumption 12. *The primary market seller chooses values p_B and p and then sets its menu of prices according to*

$$p_{Bq} = \left(\sum_j \alpha_j\right)(p_B + (1 - \psi\tau)\gamma_q) \tag{13}$$

$$p_{jq} = \alpha_j(p + \gamma_q). \tag{14}$$

Consumers who are indifferent between qualities choose the available quality with the highest value of γ_q.

The purpose of Assumption 12 is to simplify the search for the primary market seller's optimal prices by reducing it to two dimensions. The prices p_B and p are reported in counterfactual results.

Under the assumption, consumers are indifferent between all qualities and choose the best available quality. The main consequence is that the primary seller cannot intentionally create a shortage of one quality, which would push high-value consumers to buy later or consider other ticket qualities.

The only source of across-game variation in the primary seller's single-game prices is α_j. The assumption is reasonable because multiplication by α_j is the only difference between consumer values across games. Further, Assumption 12 is consistent with profit maximization in the second period if the primary seller cannot commit: the shared α_j term in consumer values and prices makes a single base price p optimal for all games.

I obtain standard errors using the bootstrap. I sample model parameters from the distribution obtained in the standard error calculations from Section 6 and run the counterfactual experiments for each set of parameters. Details can be found in the web appendix.

Resale and Partial Refunds

The primary empirical goal of the article is to compare resale to a counterfactual where the primary market seller offers partial refunds. With resale, the market is as described in Section 5.

Partial Refunds. Partial refunds are as described in Assumption 2 in Section 2, except that the primary market seller now sets a refund for each game, r_j. As before, resale is prohibited. Analogous to the resale decision rule in equation (4), consumer i requests a refund when

$$u_{ijq}(V, H_{ij}) \leq r_j. \tag{15}$$

Without a resale market, the choice set for consumers without tickets in the second period (formerly equation (9)) becomes

$$\mathcal{C}_{ij}(V, H_{ij}) = \{0, \{PM\ \ Surplus_{ijq}(V, H_{ij})\}_{q=1}^{Q}\}. \tag{16}$$

Similarly, the value of buying season tickets changes because consumers can request a refund but cannot resell. Season ticket surplus (formerly equation (8)) becomes

$$ST\ Surplus_{iq} = \sum_j \mathrm{E}_{V,H_{ij}} \Big(\max \left\{ u_{ijq}(V, H_{ij}), r_j \right\} \Big) + \delta_i - p_{Bq}. \tag{17}$$

Additionally, any refunded tickets are added back to the primary market seller's inventory and can be sold at primary market prices in the second period.

Ideally, the primary market seller would choose the profit-maximizing menu of partial refunds r_j. The optimal menu cannot be too high, or else equation (15) implies that many consumers would request refunds when the common value V is low. But refunds cannot be too low, either, because consumers need an incentive to return their tickets. Unfortunately, the data offer no guidance on what refund a consumer would accept.

In the absence of relevant data, I set the refund as 30% of the per-game price paid by season ticket buyers for zone 5 seats.[12] The refund is high enough to incentivize consumers to return their tickets, making the return rule in equation (15) plausible. They are also low enough that the model predicts few redemptions due to low values. Results in the web appendix show that the results are robust to the refund policy.[13]

Results. Table 5 shows the average performance of each counterfactual experiment over possible realizations of the common value V. Figure 7 includes plots of second-period outcomes for resale and refunds for each value of V.

The average results in Table 5 show that total welfare is maximized with partial refunds, besting resale by 0.7%, but consumer welfare is the same under both policies.[14] Profit is 2.4% higher with refunds. The magnitude of the changes should be interpreted in the context of reallocation, which only affects 7.4% of all tickets sold in the resale counterfactual.

Because refunded tickets may not be resold, 3.8% more tickets are reallocated with resale relative to the number of season tickets sold. (The number of tickets reallocated is defined as the number resold for resale, and the minimum of the number of refunds requested and primary market tickets sold for

[12] Following Assumption 12, $r_j = .3\alpha_j p_B$.

[13] Surprisingly, the robustness exercise shows that high refunds can increase the primary seller's profit by screening consumers. I limit discussion to the appendix because no ticket sellers employ such a strategy.

[14] The differences are not large relative to the standard errors. However, in all alternative parameter sets used to calculate counterfactual results for standard errors, total welfare and profit are higher with partial refunds.

refunds.) The absolute number of tickets reallocated is higher with refunds because more season tickets are sold.

The framework presented in Section 2 suggests that the change in total welfare hinges on losses to resale frictions and the benefits of flexible resale prices. Losses to frictions are present in all states and take the form of both incurred frictions and frictions that are not incurred but lead to misallocation. The benefits of flexibility are more pronounced at extreme realizations of the common value V and are shown in Figure 7.

The top-left panel confirms that resale prices vary considerably depending on the level of aggregate demand.[15] Owing to frictions, they are sold at a discount to primary market tickets in most states.

The effects of resale price variation are visible in the top-right and bottom-left panels. The top-right panel shows the number of tickets used at an average game. Although tickets sell out under both aftermarket policies when the team is good enough, more tickets are sold with resale than with refunds when the realization of V is low. When the realization of V is one standard deviation below its mean, 220 more tickets are used with resale. The difference in tickets sold, however, is small compared to the number of tickets sold in season ticket packages.

The bottom-left panel shows surplus created in the second period under each policy. When the shock V is one standard deviation above its mean, resale produces welfare 21% higher than refunds, compared to an 11% advantage at the mean shock. (Resale creates more surplus at the mean realization because fewer tickets are sold in the first period with resale.) Resale also performs better when V is one standard deviation below its mean, producing 16% more welfare than refunds.

A central reason why resale performs better when V is high is that it causes fewer tickets to be rationed when there is excess demand in the primary market. The extent of rationing is shown in the bottom-right panel. As the shock to the common value grows, demand overwhelms supply in the primary market. The probability of receiving a requested primary market ticket falls to .59 (with resale) and .66 (with refunds) when V is one standard deviation above its average. It plummets below .2 at two standard deviations. Random rationing causes consumers with relatively low values to receive the tickets; the resulting

[15]The plot also demonstrates that value changes due to shocks to V are approximately passed through completely to resale prices, an assumption used to estimate α_j, γ_q, and σ_V^2.

misallocation leads to the decline in second-period surplus for refunds in the bottom-left panel. The effect is tempered for resale because of tickets available to the highest bidder in the resale market.

Despite the greater number of tickets allocated and the more efficient outcomes at high values of V, refunds are more efficient on average. The main reasons are that (i) the extreme shocks to V do not occur often enough, and (ii) resale frictions diminish the benefits of resale.

Refunds have a clear advantage in profit, earning 2.4% more for the primary seller than resale. Interestingly, the primary seller earns more with refunds despite setting lower prices for season tickets and single-game tickets—even if expected resale revenue is subtracted from prices with resale. The reason is that it earns more from sales of single-game tickets, including by capturing some of the fees and frictions lost to resale. In fact, the primary seller earns less on season ticket sales than with refunds—despite selling more season tickets—but makes up for it by increasing primary market revenue by 62%.

Benchmark Counterfactuals

I consider several other counterfactual experiments to provide benchmarks for the performance of observed resale markets. The first two counterfactuals evaluate the effects of reallocation and primary market price rigidities. The second set of counterfactuals evaluates the effects of resale frictions, which are important to the performance of resale.

No Reallocation. I measure the overall benefits of reallocation by conducting a counterfactual without it. To eliminate reallocation, I close the resale market, as in the partial refunds counterfactual, and do not allow the primary market seller to accept returns. Any ticket sold to a consumer who receives an idiosyncratic shock goes to waste.

Flexible Prices. Many benefits of resale could be realized with refunds if primary market prices were flexible. In the flexible price counterfactual, the primary seller offers a partial refund and commits in the first period to season ticket prices and a schedule of single-game prices $p_{jq}(V)$ that varies with the realization of V.

To reduce the number of choice variables, I assume that the primary seller sets a base price p according to Assumption 12 and uses the schedule

$$p_{jq}(V) = p_{jq} + \alpha_j V. \quad (18)$$

The price schedule is nearly optimal. The adjustment for the shock V in equation (18) is essentially the same as the one the primary seller chooses in simulations where it sets a profit-maximizing price in the second period after observing V.

Resale Frictions and Fees. To quantify the importance of resale frictions to the performance of resale, I conduct counterfactuals with resale markets but no resale frictions ($\lambda_s = 0$) and neither resale frictions nor fees ($\lambda_s = 0$ and $\tau = 0$).

Results. Results are in Table 6. The results of the counterfactual without reallocation confirm that there are substantial benefits to both aftermarket policies. Without reallocation, total welfare is 4.7% lower than with resale, consumer surplus is 7% lower, and profit is 2.4% lower. The results reinforce that primary market sellers can benefit from reallocation.

The results with flexible primary market prices suggest that a partial refund scheme would easily be optimal if there were no price rigidities in the primary market. Total welfare is 3% higher than in the main refund counterfactual, consumer welfare is 2.5% higher, and profit is 3.3% higher. A notable difference compared to the refund counterfactual with price rigidities is that fewer season ticket packages are sold and single-game tickets are generally cheaper, yet tickets sell out for all realizations of aggregate shocks. One reason for the difference is that the primary market seller can charge consumers for their expected values in the first period, but second-period outcomes are volatile when prices are rigid. With flexible prices, the primary market seller has less incentive to shift sales to the first period.

The rightmost columns measure the effect of resale frictions. Without resale frictions, total welfare would be 0.9% higher than with refunds, and consumer welfare would be unchanged.

Removing fees in addition to resale frictions does not cause the primary market seller to change the allocation of tickets—the only effect is to transfer surpus from the resale market operator to the primary market seller. The primary market seller does so by raising its season ticket prices, charging consumers for the resale revenue gained when the fee disappears. The change in fees is enough to change the optimal strategy. The primary market seller earns 1.2% more profit with frictionless resale and no fees than with refunds, but

1.5% less profit with fees. The effects of removing fees would be different in a model with an intensive margin to resale.

Discussion. Although brokers are not expected to be important in the empirical setting, the results shed some light on their effects. As in settings with systematic underpricing and brokers, demand outstrips supply in the primary market when the realized aggregate shock is high. The analogy is not perfect. Underpricing is known at the start of the market in settings with intentional resale and brokers. In contrast, underpricing is only known after shocks are realized in this article and so the initial allocation is not affected. Moreover, consumers cannot buy underpriced tickets to resell within the second period.

The results are still illustrative for a central question regarding brokers: without brokers (or resale), how severe would misallocation be? Figure 7 suggests that misallocation from primary market rationing is significant enough that resale, and brokers, could be valuable. As aggregate shocks grow, the welfare produced by resale increases, but misallocation from rationing causes welfare to decline with refunds. Brokers may reduce rationing and thus raise welfare in such cases. However, the model does not capture several harms of brokers.

One potential harm is that brokers could have market power. Another is that there are costs to acquiring tickets in settings where they are underpriced in the primary market. The resulting costs—an important feature in Leslie and Sorensen (2014)—reduce the benefits of brokers but cannot be measured in this setting. If, contrary to expectations, brokers are important in the market, the model would overstate the rate of idiosyncratic shocks and, if brokers extract more surplus than casual resellers, would also overstate consumer surplus.

The comparison between resale and refunds may change if primary and resale markets are sold in a single integrated market. The promise of integrated resale is to reduce the significant frictions associated with resale. But there are risks. An integrated reseller may acquire market power in the resale market that lets it raise fees or gives it an incentive to distort the allocation of tickets. The model is not able to predict the effects of integrated resale because the change in frictions is unknown and there is no intensive margin to resale in the model, which would affect the optimal resale fee.

The counterfactual results also inspire a practical concern: why are refunds rare in ticket markets? After all, events like concerts have little demand uncertainty, which removes the benefits of resale studied in this article. A likely reason is that restrictions on resale and ticket transfers are unpopular, leaving unlucky consumers with expensive tickets they cannot use (see e.g. Pender (2017)). Enacting a refund policy would also require the primary market seller to consider market dynamics: refunding a ticket a week before the show is more valuable than refunding it an hour before. The challenge may discourage event organizers from abandoning familiar resale policies. Finally, although the artists who set prices too low may not be trying to maximize welfare, the results suggest that the advantages of refunds may be reduced when tickets are underpriced.

8 Conclusion

When consumers receive stochastic demand shocks, the initial allocation of goods can be suboptimal. Society can benefit from aftermarket policies that cope with shocks, but it is unclear which policy is best. I show that the optimal aftermarket policy depends on the relative importance of aggregate shocks and resale frictions; I then estimate a structural model describing the salient shocks and frictions in the market for college football tickets and evaluate each policy in counterfactual experiments.

The results suggest that refunds are more efficient than the status quo of resale. In counterfactual experiments, total welfare is 0.7% higher with refunds, profit is 2.4% higher, and consumer welfare does not change. The differences in welfare and profit are meaningful given that 7.4% of tickets are reallocated in equilibrium.

However, the average performances of the policies mask important differences. Resale performs relatively better when there are more extreme aggregate shocks, allocating more tickets when the shock is low and avoiding rationing when the shock is high. Such conclusions are only possible in a model with aggregate demand shocks.

The article has three core implications for our understanding of resale and aftermarkets. First, the framework demonstrates that resale can be valuable in markets with primary market rigidities, aggregate uncertainty, and low resale frictions. The market for college football tickets includes both rigidities and aggregate uncertainty, but resale frictions are significant enough for refunds to be optimal. In similar markets without primary market rigidities, like airlines and hotels, refunds are a natural choice.

Second, the comparison between resale and refunds informs how to run aftermarkets. The results imply that refund strategies can be superior in perishable goods markets even when there is significant aggregate demand uncertainty. A driver of the benefits is the removal of frictions associated with resale.

Third, the article provides empirical evidence on the effects of resale. Whether primary market sellers of perishable goods profit from resale is ambiguous in theory, and this article shows that primary sellers benefit in practice: resale raises profit by 2.4% compared to not reallocating. The effects of not reallocating on welfare inform policy on ticket resale. Total and consumer welfare fall significantly, by 4.7% and 7% compared to resale. Society would benefit from a legal right to resell tickets in cases where the primary seller does not offer an alternative method of reallocation.

References

Bhave, Aditya and Eric Budish. "Primary-Market Auctions for Event Tickets: Eliminating the Rents of 'Bob the Broker'?" *American Economic Journal: Microeconomics*, 15, 142–170 (2023).

Board, Simon and Andrzej Skrzypacz. "Revenue Management with Forward-Looking Buyers." *Journal of Political Economy*, 124, 1046–1087 (2016).

Cachon, Gérard P. and Pnina Feldman (2021), "Why Facilitate Ticket Exchanges and Other Reselling Markets: Pricing Capacity and Recourse Strategies." Technical report, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3235757.

Chen, Jiawei, Susanna Esteban, and Matthew Shum. "When Do Secondary Markets Harm Firms?" *American Economic Review*, 103, 2911–2934 (2013).

Collard-Wexler, Allan. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica*, 81, 1003–1037 (2013).

Courty, Pascal. "Ticket Pricing Under Demand Uncertainty." *Journal of Law and Economics*, 46, 627–652 (2003).

Courty, Pascal. "Ticket resale, bots, and the fair price ticketing curse." *Journal of Cultural Economics*, 43, 345–363 (2019).

Cui, Yao, Izak Duenyas, and Özge Şahin. "Should Event Organizers Prevent Resale of Tickets?" *Management Science*, 60, 2160–2179 (2014).

De Avila, Joseph. "Miami's Messi Mania: Ticket Prices Soar, Bars Are Packed." *The Wall Street Journal*, URL https://www.wsj.com/articles/lionel-messi-inter-miami-mls-soccer-south-florida-ticket-prices-sports-bars-21f52ead (2023).

Dilmé, Francesc and Fei Li. "Revenue Management without Commitment: Dynamic Pricing and Periodic Flash Sales." *Review of Economic Studies*, 86, 1999–2034 (2019).

Egelko, Bob. "Season ticket holder sues 49ers over new resale policy." *San Francisco Chronicle*, URL https://www.sfgate.com/49ers/article/Season-ticket-holder-sues-49ers-over-new-resale-6438554.php (2015).

Goldberg, Brett (2019), "Stubhub fees — the truth about buyer and seller fees." https://www.tickpick.com/blog/stubhub-buyer-seller-fees/. Accessed March 19, 2021.

Hendel, Igal and Alessandro Lizzeri. "The Role of Leasing under Adverse Selection." *Journal of Political Economy*, 110, 113–143 (2002).

Hodgson, Charles. "Trade-ins an Transaction Costs in the Market for Used Business Jets." *American Economic Journal: Microeconomics*, 15, 350–391 (2023).

Jeon, Jihye. "Learning and Investment under Demand Uncertainty in Container Shipping." *RAND Journal of Economics*, 51, 226–259 (2022).

Kalouptsidi, Myrto. "Time to Build and Fluctuations in Bulk Shipping." *American Economic Review*, 104, 564–608 (2014).

Lazarev, John (2013), "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets." Working paper, New York University.

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

Lewis, Michael, Yanwen Wang, and Chunhua Wu. "Season Ticket Buyer Value and Secondary Market Options." *Marketing Science*, 38, 973–993 (2019).

Pender, Kathleen. "Think carefully before buying 'paperless' tickets that require credit card entry." *San Francisco Chronicle*, URL https://www.sfchronicle.com/business/networth/article/Think-carefully-before-buying-paperless-11186609.php (2017).

Satariano, Adam. "The Case of the Stubbed Hub." *Bloomberg Businessweek*, URL https://www.bloomberg.com/news/articles/2015-02-19/stubhub-faces-pressure-from-ticketmaster-and-ebay-its-own-parent (2015).

Squire Patton Boggs LLP (2017), "Secondary Ticket Marketplace Guide to US Ticket Resale Regulations." Report, Squire Patton Boggs LLP.

StubHub (2021), "Stubhub's ticket fees." https://support.stubhub.com/en/support/solutions/articles/61000276392-stubhub-s-ticket-fees. Accessed March 19, 2021.

Sweeting, Andrew. "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets." *Journal of Political Economy*, 120, 1133–1172 (2012).

Waisman, Caio. "Selling Mechanisms for Perishable Goods: An Empirical Analysis of an Online Resale Market for Event Tickets." *Quantitative Marketing and Economics*, 19, 127–178 (2021).

Williams, Kevin R. "The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets." *Econometrica*, 90, 831–858 (2022).

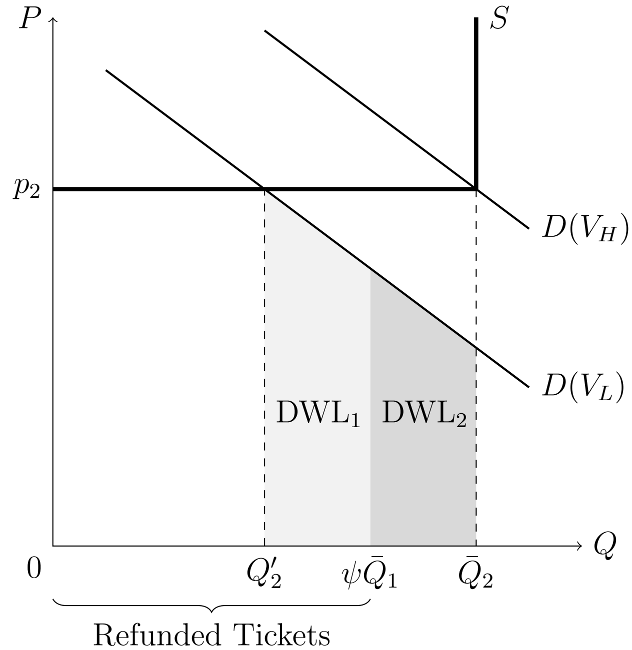

Figure 1: The second period of a market with refunds and uncertain demand.

Table 1: Primary market single-game and season ticket prices. Table excludes the canceled game. Season ticket prices are prorated to reflect the cancellation.

Game	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
1	70	60	50	40	30
2	70	60	55	45	30
3	70	60	50	40	30
4	70	60	55	45	30
5	60	55	40	35	30
Season Tickets	315	270	216	179	125
Face Value Sum	340	295	250	205	150

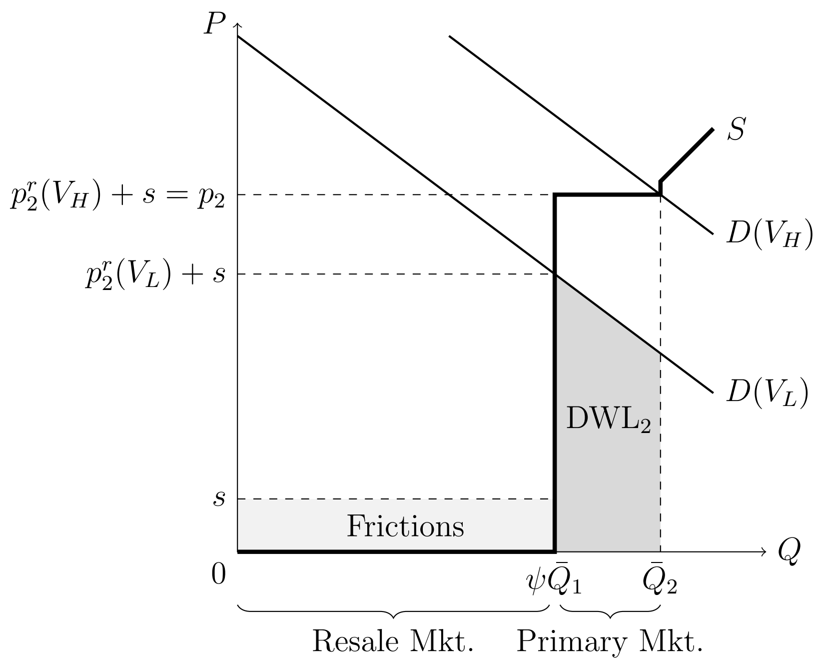

Figure 2: The second period of a market with resale and uncertain demand.

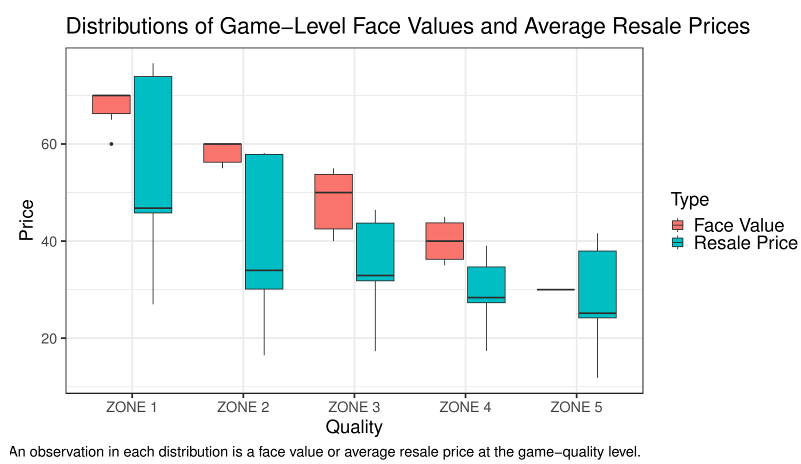

Figure 3: Distributions of mean fee-inclusive per-game resale prices and face value. Based on university policies and StubHub data.

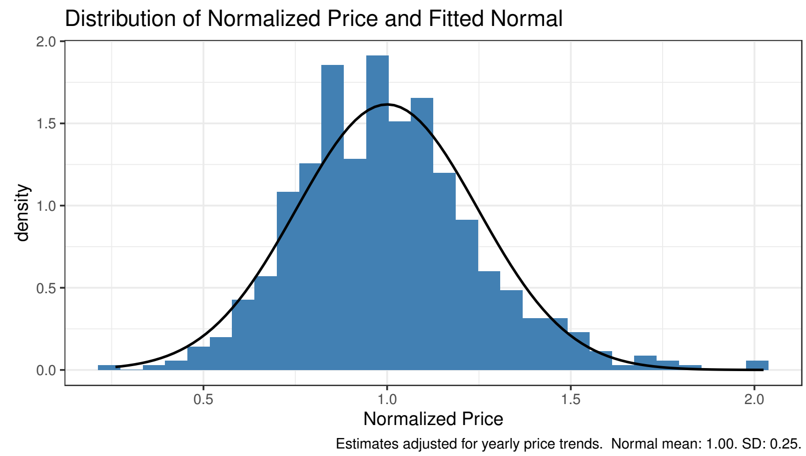

Figure 4: Distribution of resale prices normalized by sample university means, adjusted for yearly trends, and fitted normal distribution. Based on SeatGeek data.

Parameter Description	Notation	Estimate	Std. Err.
Resale Fee (%)	τ	0.22	-
Idiosyncratic Shock Rate	ψ	0.08	-
Preference for Game 1	α_1	1.00	-
Preference for Game 2	α_2	1.67	(0.032)
Preference for Game 3	α_3	1.01	(0.023)
Preference for Game 4	α_4	1.60	(0.029)
Preference for Game 5	α_5	0.56	(0.015)
Preference for Quality 1	γ_1	21.95	(0.689)
Preference for Quality 2	γ_2	9.90	(0.611)
Preference for Quality 3	γ_3	4.37	(0.569)
Preference for Quality 4	γ_4	-0.70	(0.619)
Preference for Quality 5	γ_5	0.00	-
SD of Common Value	σ_V	8.08	(0.293)

Table 2: Estimates for parameters that do not require model simulations. Standard errors calculated using the bootstrap.

Table 3: Estimated parameters from the second stage.

Parameter Description	Notation	Estimate	Standard Error
Mean Resale Friction	λ_s	78.05	(0.68)
Mean Consumer Type	λ_ν	17.60	(0.03)
High-Type ST Benefits	δ_H	5.60	(0.04)
Low-Type ST Benefits*	δ_L	-203.06	(1.67)
Pct. High-Type ST Benefits	ζ	0.81	(0.004)

* Estimate and standards errors are for the upper bound of the identified set.

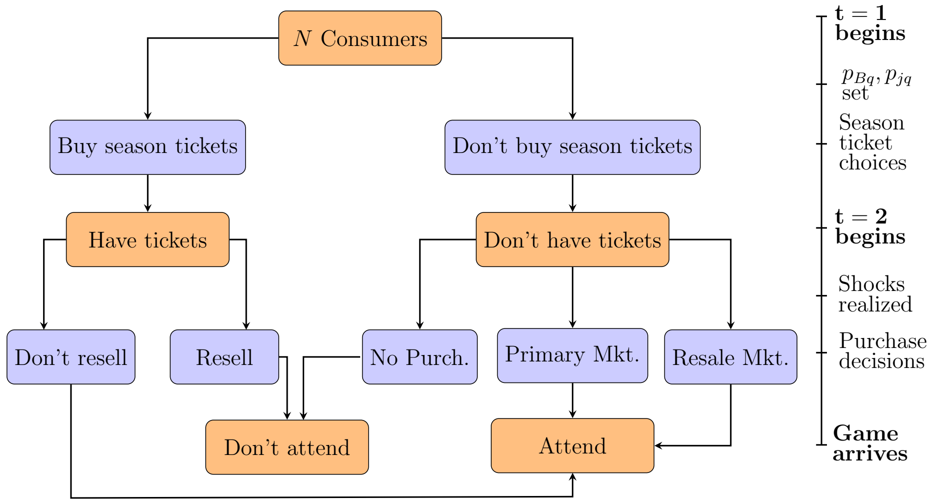

Figure 5: Model timeline and outline for consumer arrivals and choices, where the second period is shown for a single game j. Decisions are shown in blue.

Table 4: Observed and model-implied quantities of season tickets.

Moment	Model-Implied	Observed
Season Tickets Sold	24543	22471

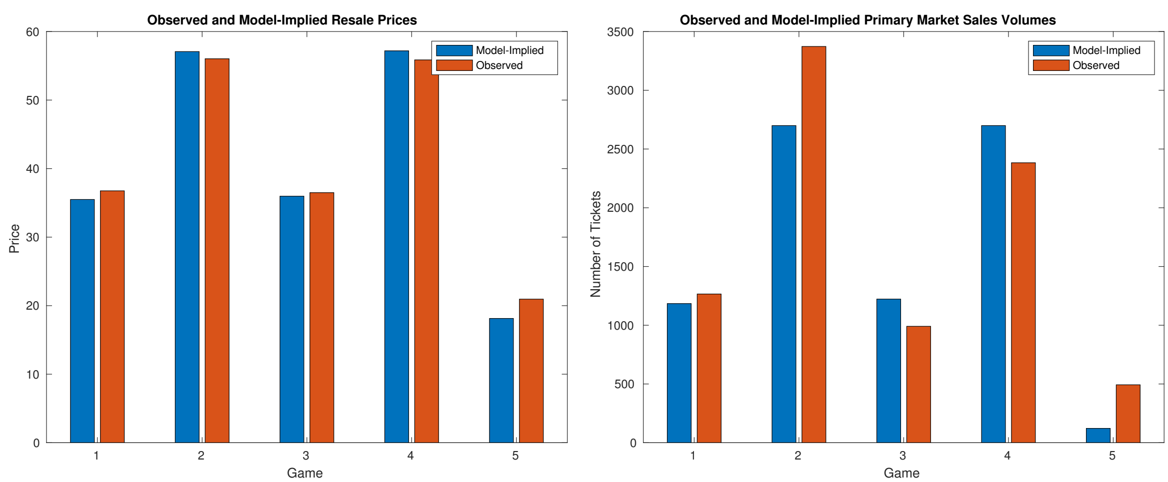

(a) Resale prices (b) Primary market quantities

Figure 6: Observed and model-implied resale prices and primary market quantities for the realized value of V for each game.

	Resale	Refunds
Total Welfare (mn)	\$10.11	\$10.18
	(0.12)	(0.12)
Profit (mn)	\$7.17	\$7.34
	(0.09)	(0.09)
Consumer Welfare (mn)	\$2.84	\$2.84
	(0.04)	(0.04)
Resale Fees (mn)	\$0.10	\$0.00
	(0.00)	(0.00)
Tickets Resold or Refunded (1000)	10.28	10.78
Reallocated Tickets (1000)	10.28	10.35
Season Ticket Buyers (1000)	25.83	27.03
Season Ticket Base Price	\$31.82	\$30.71
Single Game Base Price	\$42.22	\$40.98

Table 5: Average counterfactual results across realizations of V for resale and refunds. Standard errors calculated using the bootstrap and shown in parentheses.

	Resale	Flex. Prices	No Reall.	$\lambda_s = 0$	$\lambda_s = \tau = 0$
Total Welfare (mn)	\$10.11	\$10.49	\$9.64	\$10.27	\$10.27
	(0.12)	(0.13)	(0.12)	(0.22)	(0.20)
Profit (mn)	\$7.17	\$7.58	\$7.00	\$7.23	\$7.43
	(0.09)	(0.09)	(0.09)	(0.13)	(0.13)
Consumer Welfare (mn)	\$2.84	\$2.91	\$2.64	\$2.84	\$2.84
	(0.04)	(0.04)	(0.04)	(0.09)	(0.08)
Resale Fees (mn)	\$0.10	\$0.00	\$0.00	\$0.20	\$0.00
	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)
Season Ticket Buyers (1000)	25.83	26.42	24.96	27.18	27.18
Season Ticket Base Price	\$31.82	\$31.17	\$31.22	\$33.00	\$37.15
Single Game Base Price	\$42.22	\$37.40	\$41.47	\$41.08	\$41.08

Table 6: Average counterfactual results across realizations of V for experiments with flexible prices, no reallocation, and no resale frictions. Standard errors calulated using the bootstrap and shown in parentheses.

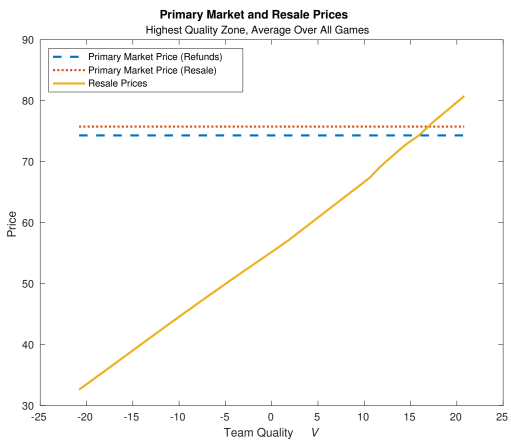

(a) Primary and resale market prices.

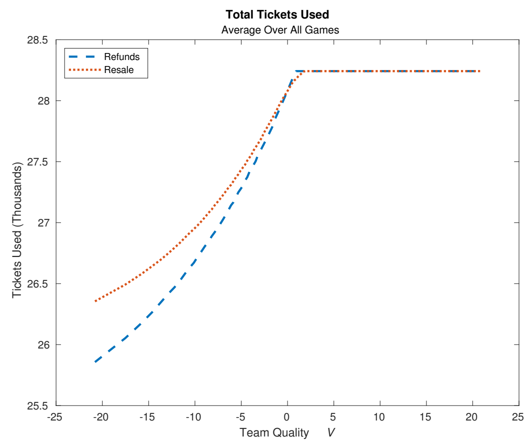

(b) Number of tickets used, average across games.

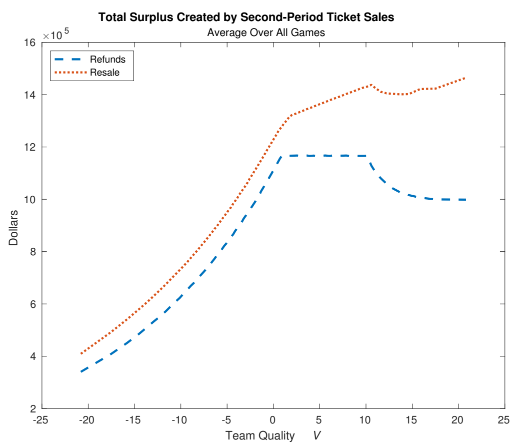

(c) Total surplus created from period two sales.

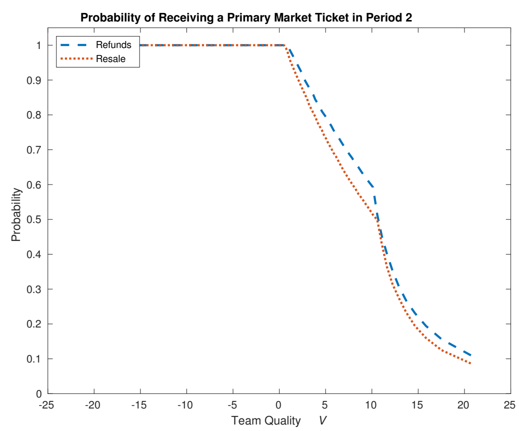

(d) Primary market rationing probability.

Figure 7: Market outcomes as the common value V varies.