

Price Salience and Product Choice*

Tom Blake[†]
eBay Research

Sarah Moshary[‡]
University of Chicago
Booth School of Business

Kane Sweeney[§]
Uber

Steve Tadelis[¶]
UC Berkeley,
NBER & CEPR

July 7, 2020

Abstract

Online vendors often employ drip-pricing strategies, where mandatory fees are displayed at a later stage in the purchase process than base prices. We analyze a large-scale field experiment on StubHub.com and show that disclosing fees upfront reduces both the quantity and quality of purchases. The effect of salience on quality accounts for at least 28% of the overall revenue decline. Detailed click-stream data shows that price shrouding makes price comparisons difficult and results in consumers spending more than they would otherwise. We also find that sellers respond to increased price obfuscation by listing higher quality tickets.

*We are grateful to executives and employees at Stubhub for sharing the data for this study. We thank numerous seminar participants for helpful comments.

[†]tomblake@gmail.com

[‡]sarah.moshary@chicagobooth.edu. Address: 5807 S Woodlawn Avenue, Chicago, IL 60637

[§]kane@uber.com

[¶]stadelis@berkeley.edu

1 Introduction

The past two decades have witnessed a steady shift in purchasing from brick-and-mortar stores to online retailers and marketplaces. A common pricing strategy used by online vendors—most notably for event ticket sales—is “drip pricing,” where mandatory fees are disclosed at a later stage in the consumer’s purchasing process than the base price of a good. Textbook models of consumer choice assume that economic agents are rational and sophisticated in their ability to discern a product’s true price, implying that purchase decisions fully account for any fees, taxes, or add-on features. However, a growing literature demonstrates that consumers often struggle to determine final prices. For example, Chetty et al. (2009) document that tax salience affects consumers’ decisions to purchase personal care goods in grocery stores, implying that consumers have trouble inferring final prices when taxes are not displayed on the shelf. Morowitz et al. (1998) find that students in a lab react less to surcharges presented as percentages rather than dollars, suggesting a cognitive difficulty in calculating prices. Hossain and Morgan (2006) and Brown et al. (2010) present evidence that eBay buyers respond more to list price than to shipping cost.

Studies have therefore demonstrated that consumers are more likely to purchase goods when fees are obfuscated. Our paper contributes in two ways. First, we employ a large-scale field experiment involving millions of online consumers to confirm what small-scale studies have shown, and we use our detailed data to expose behaviors along the purchase funnel. Second, and more novel, we show that price salience affects not only whether a consumer chooses to purchase *any* product, but also affects their choice of *which* product to purchase. Our setting is a secondary marketplace for event tickets where more expensive tickets are associated with better (higher quality) seats. We show that when fees are less salient consumers are more likely to select and purchase more expensive tickets. Intuitively, reducing the salience of a percent-based purchasing fee makes all goods appear less expensive, enticing more consumers to select and then purchase a ticket. Because a percentage fee levies a larger fee level for more expensive goods, salience also changes the perceived marginal cost of quality. As a result, reducing salience encourages consumers to substitute to high quality tickets. We therefore offer a more complete analysis of the effect of price salience on consumer choice first by demonstrating effects on the intensive margin, and second, by quantifying the relative importance of both the extensive and intensive margins in our setting.¹

We begin our analysis by presenting two hypotheses that follow from the existing theoretical literature: first, that consumers are more likely to purchase goods if fees are obfuscated; and second, that consumers are more likely to purchase expensive, high-quality goods if fees are obfuscated. The former effect has been documented by many studies, but the latter has not

¹In their working paper version, Chetty et al. (2009) note that the revenue effect is bigger than the quantity effect, which is potentially due to consumers switching to lower priced items. Their data is insufficient to investigate that possibility further.

been explored because of data limitations in earlier work.

We take these predictions to data generated from a large-scale field experiment conducted by StubHub, a leading online secondary ticket marketplace. Before the experiment was launched in August 2015, the platform used an Upfront Fee (UF) strategy, where the site showed consumers the final price including fees and taxes from their very first viewing of ticket inventory. The platform then experimented with a Back-end Fee (BF) strategy, where mandatory fees were shown only *after* consumers had selected a particular ticket and proceeded to the checkout page.


StubHub randomly selected 50% of U.S. users for the BF experience, while the remaining 50% were assigned to the UF experience. The experiment provides exogenous variation in fee salience in a setting with rich data on consumer choices, including choice sets, signals of purchase intent (e.g. product selection and clicks towards checkout), and final purchases. These rich data allow us to infer the effect of salience on both the extensive and intensive margins of product choice. Our empirical results support our hypotheses: price obfuscation distorts both quality and quantity decisions. A simple lower-bound estimate shows that the intensive margin—how expensive of a ticket to buy—accounts for at least 28% of the increase in revenue raised from Back-end Fees.

Further analysis of detailed individual-level clickstream data suggests that Back-end Fees play on consumer misinformation. UF users are more likely to exit before exploring any ticket, while BF users differentially exit at checkout, when they first see the fee. Furthermore, BF users go back to examine other listings more often than their UF counterparts. They are more likely to go back multiple times, which suggests that Back-end Fees make price comparisons difficult. Finally, Back-end Fees affect even experienced users, although on a smaller scale, which is consistent with consumers facing optimization costs even when they anticipate a fee, as in Morowitz et al. (1998).

We also investigate how sellers who list on StubHub respond to the change in fee salience on the platform following the experiment’s conclusion, when StubHub shifted the whole site to Back-end Fees. Because Back-end Fees cause buyers to purchase more tickets, and in particular more expensive tickets, the two-sided nature of the platform should incentivize sellers to list relatively more expensive, high quality tickets. Using row-numbers as a proxy for quality, our analysis shows that sellers indeed choose to list higher quality tickets after the transition to Back-end Fees. We also find that sellers respond in how they set prices; in particular, they are more likely to set list prices at round numbers. Hence, consistent with Ellison and Ellison (2009), we find that sellers respond to the change in buyer experience.

As a robustness check, we present evidence on price salience from an earlier experiment at StubHub performed in 2012. One advantage of this earlier experiment is that StubHub’s default user experience during the experiment was BF, as shown in figure I. Thus, comparing the results from the 2012 and 2015 experiments can shed light on whether the effect of

Figure I: Timeline of Fee Presentation at StubHub



salience depends on the initial environment. Our findings indicate that the effect of salience is remarkably similar across the two experiments. A second feature of the 2012 experiment is that it randomized fee presentation across events, rather than across users. This experiment design circumvents interference from device-switching, when a user is randomized into different conditions on their mobile/laptop/desktop computers. Reassuringly, the results are broadly consistent with our findings from the 2015 experiment, indicating that this concern is not first-order in our setting.

Our paper also contributes to studies of alternative methods of obfuscation, such as add-on pricing and partitioned pricing. Ellison (2005) and Gabaix and Laibson (2006) explore models where some consumers ignore the price of complimentary goods (e.g., parking at a hotel) when making purchase decisions. Predictions from these models have been examined in recent empirical work, such as Ellison and Ellison (2009) and Seim et al. (2017) (see Heidhues and Köszegi (2018) for an overview). In the language of Gabaix and Laibson (2006), StubHub fees constitute surcharges rather than add-ons because they are unavoidable. We might interpret the StubHub fee as a form of partitioned pricing because it is broken out from the base price of the ticket (see Greenleaf et al. (2016) for a review of the partitioned pricing literature). One interpretation of our findings is that salience amplifies the effect of partitioned pricing. Salience may therefore help explain the persistence of markups and price dispersion in online markets, as documented by Brynjolfsson and Smith (2001), among others.

Closest to our paper is a recent study by Dertwinkel-Kalt et al. (2019), who examine the online purchase behavior of over 34,000 consumers of a large German cinema that obfuscated a surcharge for 3D movies until check-out. They find that consumers initiate a purchase process more often when surcharges are obfuscated, but they also drop out more often when the overall price is revealed at check-out. In their setting, these two effects counteract each other so that the demand distribution is independent of the price presentation. Hence, our findings differ from theirs in three important ways. First, as in previous studies, we find that obfuscation increases demand, meaning that the increased rate of purchase initiation outweighs the increased drop-out rate caused by obfuscation. Second, our richer setting allows us to document how salience affects the intensive margin. Third, and most importantly, our findings contravene the argument in Dertwinkel-Kalt et al. (2019) that the salience effects

documented in previous studies, such as Chetty et al. (2009), Taubinsky and Rees-Jones (2018) or Feldman and Ruffle (2015), do not generalize to online settings because e-commerce transactions often involve a single, focal product. Dertwinkel-Kalt et al. (2019) argue further that low cancellation costs, such as clicking back on a page, limit the effectiveness of practices like drip-pricing. Our results suggest otherwise, as we find a large effect of price salience in a large online marketplace with very low cancellation costs.

The next section presents a standard framework for consumer choice with price obfuscation and describes its empirical implications. Section 3 discusses the experiment run at StubHub, as well as the data used in the analysis. Section 4 describes robustness checks on the randomization, while section 5 presents our main results. Section 6 contains evidence on mechanisms and section 7 explores two-sided market responses. Section 8 concludes.

2 Consumer Choice with Fee Obfuscation: Hypotheses

As a starting point, we build on the insights of Bordalo et al. (2013) and DellaVigna (2009), who each present simple models of consumer choice that explore the impact of price salience on purchase decisions. In the Appendix we present a simple model based on these studies that formalizes our two main hypotheses: that obfuscating check-out fees causes more consumers to purchase goods, and that the goods they purchase will be more expensive and of higher quality compared to an environment with upfront fees.

In our setting, consumers visit the Stubhub website—a platform for secondary market ticket sales—in order to purchase tickets for events. As we describe in more detail in Section 3, final prices of tickets are made up of two components: a list price set by sellers, and fees set by Stubhub. We consider two salience conditions under which consumers make purchase decisions: the first is the “Upfront Fee” (UF) condition, where the the final purchase price including all fees is shown to consumers upfront when they search for available tickets, and the second is the “Back-end Fee” (BF) condition, where consumers observe only list prices set by sellers when searching for tickets and the fees imposed by Stubhub are revealed only after the consumer proceeds to the checkout stage with a particular ticket. Section 3 offers more details about the experiment’s design and execution.

Consider the UF case. If all ticket prices exceed a consumer’s willingness to pay, then she will not buy any ticket. If some are priced below her willingness to pay, then she will buy the ticket that maximizes her net surplus. Naturally, the higher her value for a given event, the more likely she is to purchase a ticket. Conditional on purchasing, the more she values the event, the more likely she is to buy an expensive, high quality. Finally, because fees are included upfront, the purchase price the consumer faces at checkout is identical to the price that she saw on the listing page.

Now consider the BF case, where fees are revealed for the first time at checkout. Because

fees amount to about 15% of the list price, if a consumer considers only the list price, then all tickets appear to be 15% cheaper during the consumer’s search phase. The consumer therefore makes a choice from a seemingly cheaper set of tickets. This is akin to reducing the salience of prices relative to quality as in Bordalo et al. (2013) and is also similar to the way Finkelstein (2009) models salience. As a consequence, consumers who would not have chosen any ticket under UF may believe that they have found a cheap enough ticket under BF to warrant purchase, and proceed to the checkout page with that ticket in hand. Upon reaching the checkout and purchase page, the ticket’s *actual* price—including all fees—is revealed. Absent behavioral biases, the consumer ought to exit without buying the ticket, but we assume that some consumers will complete their purchase due to loss aversion or other behavioral biases.² This results in the well established and previously tested hypothesis:

(1) Quantity Effect: *A consumer is more likely to purchase under BF than under UF.*

One of our main innovations compared to the previous literature is going beyond this quantity effect to explore how the *composition* of products purchased changes across the two conditions. To see this, consider a consumer who would have chosen a ticket listed at \$100 under UF. Under BF, she instead selects a \$100 ticket to which a \$15 fee will be added at checkout, so that her purchase under BF is equivalent to a \$115 ticket in the UF condition. With no behavioral biases and no search costs, this BF consumer would go back to the listing page and select a ticket that maximizes her utility (an \$87 ticket, that will cost just about \$100 after the fee is included at checkout). We again assume that some consumers will not re-optimize and instead will purchase their initial choice due to loss aversion or search costs, resulting in the following hypothesis that has not been analyzed previously in the literature:

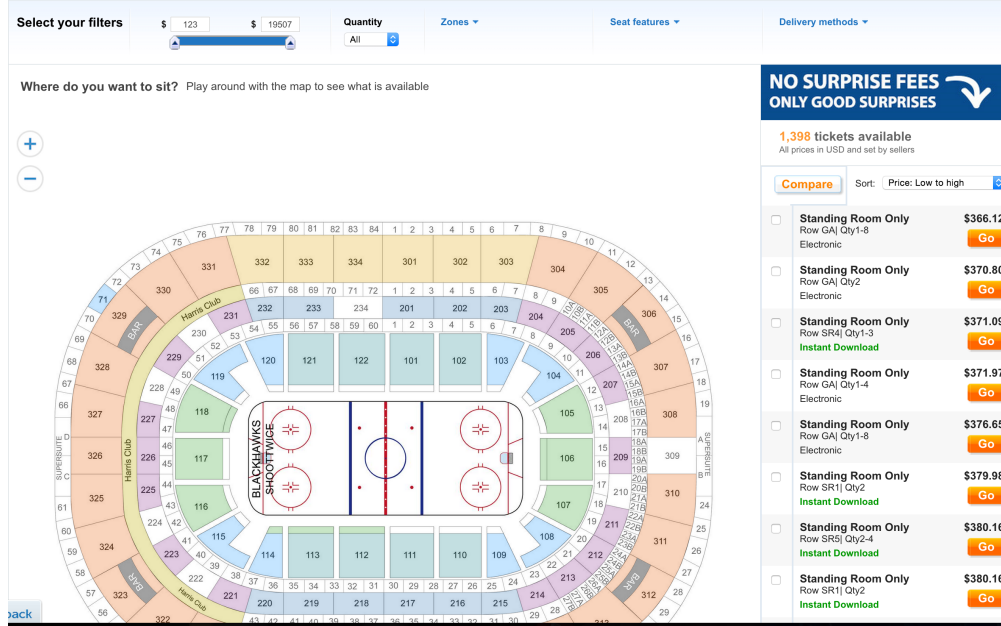
(2) Quality Upgrade Effect: Consumers who buy tickets under both UF and BF conditions will purchase higher quality and more expensive tickets under BF.

The earlier salience literature overlooks this effect, perhaps because previously studied settings offered little to no vertical product differentiation (e.g., shipping fees as in Brown et al. (2010), Electronic Toll Collection systems as in Finkelstein (2009) or supermarket beauty aids as in Chetty et al. (2009)). Indeed, the log-log demand specification favored by earlier work leaves no scope for quality upgrades.

The Quality Upgrade Effect emphasizes how identification strategies must respect the impact of salience on quality choice. Consider the alcohol sales analysis of Chetty et al. (2009). They compare an excise (lump sum) tax to a sales (percentage) tax. The excise tax should arguably have no effect on the quality of beer chosen (conditional on purchase), since

²An alternative explanation is that by entering payment information en route to the checkout page, BF users face lower barriers to purchase than UF users. We find this explanation unlikely because hassle costs must be very large to explain the salience effects.

Figure II: Event Page (UF users)



it makes each can of beer “in the choice set” more expensive by the same amount. The sales tax, however, may affect both the quantity and quality margins, since it is a percentage of the price. Simple comparisons of the revenue effects of excise and sales tax salience may therefore lead to inconclusive results.

The next section described the experiment in detail and elaborates our empirical strategy for separately estimating the quantity effect, bounds on the quality upgrade effect, revenue effects, and the change in the average purchase price.

3 Experimental Design

We exploit an experiment in price salience performed on StubHub, a platform for secondary market ticket sales. Between January 2014 and August 2015, the platform showed all fees upfront, so the initial prices a consumer saw when browsing ticket inventory was the final checkout price. Figure II shows an Event page, which is what consumers see when they select an event that they are interested in attending. Ticket inventory is listed on the right, and prices including all fees are presented for each ticket.

Between August 19 and August 31 of 2015, the firm ran an experiment where treated consumers were initially shown ticket prices without fees.³ For treated customers, fees were added at the checkout page, much like sales taxes at the register of a store. We refer to this user experience as Back-end Fees.⁴ StubHub’s fee structure is non-linear: the buyer fee is

³Ethan Smith, “StubHub Gets Out of All-In Pricing.” *Wall Street Journal*, August 31, 2015.

⁴Ticketmaster and other platforms also employ a similar Back-End Fee pricing scheme.

Figure III: Treatment versus Control Experiences

Upfront Fees (Event Page)

Back-end Fees (Event Page)

Both lead to the same checkout:

15% of the ticket price plus shipping and handling, if applicable. StubHub also charges seller fees, which peak at 15%.

The experimental condition was assigned at the cookie-level, which identifies a browser on a computer. Half of U.S. site visitors were assigned to the treatment (BF) group at their first touch of an event page. On the event page, users are shown a list of tickets. Consumers assigned to the pre-experimental UF experience (the control group) were shown conspicuous onsite announcements confirming that the prices they saw upfront included all charges and fees. On the other hand, treated users in the BF group were shown only the base price when they perused available listings. Once a user in the BF group selected a ticket, they were taken to a ticket details page, where they could log in to purchase the ticket and then review the purchase. It is at this point that the BF group was shown the total price (ticket cost plus fees and shipping charges). Users could then checkout or abandon the purchase. Figure III shows the different prices on the Event page that result in the same price on the checkout page for treatment and control.

First, we exploit the randomization to estimate the quantity effect described in section 2 as the difference in purchase probabilities between UF and BF users.⁵ Because sellers on StubHub cannot price discriminate between BF and UF users, we need not worry that the two groups face different prices because of the treatment (nor do we include other control variables). In practice, we estimate the following regression equation via OLS, where Q_i is an

⁵Using the potential outcomes notation, we can write the quantity effect as $\Delta Q = E[Q_i|T_i = 1] - E[Q_i|T_i = 0]$.

indicator that consumer i purchases a ticket and T_i is a BF treatment indicator:

$$Q_i = \alpha + \beta T_i + \epsilon_i. \quad (1)$$

The parameter β represents the difference in the levels of purchasing (Q_i) for BF compared to UF users. To protect business-sensitive information, however, we report estimates of $\frac{\beta}{\alpha}$, which is the percent change in the likelihood of purchase for BF users.

Measuring the Quality Upgrade Effect is challenging because the random assignment of the BF experience changes the identity of the marginal consumer. Our intuition, developed more fully in Appendix A, suggests that the marginal consumer who purchases under BF has a lower valuation for the event and chooses lower quality tickets.⁶ Measuring the Quality Upgrade Effect requires adjusting for this selection. Namely, conditional on i making a purchase, let P_i be the purchase price of the ticket that i selects. Let Q_{i0} be an indicator for whether consumer i purchases a ticket when he observes fees upfront ($T_i = 0$) and Q_{i1} for when he observes fees at the back-end ($T_i = 1$). We formulate the Quality Upgrade Effect using the potential outcomes notation as:

$$\text{QUE} = E[P_i|Q_{i0} = 1, T_i = 1] - E[P_i|Q_{i0} = 1, T_i = 0] \quad (2)$$

The second term is observed by the econometrician and is the average price of tickets purchased by UF users. The challenge is that the econometrician cannot observe the first term, which is the average price of tickets UF users would buy *if they were exposed* to the BF treatment. Instead, we observe the change in the average price, conditional on purchasing:

$$\begin{aligned} \Delta P &= E[P_i|Q_{i1} = 1, T_i = 1] - E[P_i|Q_{i0} = 1, T_i = 0] \\ &= \text{QUE} + E[P_i|Q_{i1} = 1, T_i = 1] - E[P_i|Q_{i0} = 1, T_i = 1] \\ &\leq \text{QUE}. \end{aligned} \quad (3)$$

Equation 3 shows that the change in the average purchase price (ΔP) combines two separate effects: first, the Quality Upgrade Effect, where BF encourages consumers to purchase more expensive tickets than they would otherwise, and second, a change in the marginal consumer, as BF induces more consumers to purchase tickets.⁷ The former increases the average purchase price while the latter depresses it (because marginal consumers buy cheaper tickets). We therefore use ΔP as a lower bound for the Quality Upgrade Effect; we estimate (3) using regression specification (1) with price as the left-hand side variable.

We note that the change in average purchase price is inherently interesting in this setting, as it maps to a change in platform revenue. We decompose the change in revenue from

⁶In the language of the model that appears in the appendix, the marginal consumer has a lower θ .

⁷The derivation employs the standard monotonicity of choice for a given consumer (i.e., $Pr\{Q_{i1} = 1|Q_{i0} = 1\} = 1$).

treatment as⁸

$$\Delta E[R_i] = \underbrace{\Delta E[P_i|Q_i = 1]}_{\Delta P} \cdot E[Q_i] + \underbrace{\Delta E[Q_i]}_{\Delta Q} \cdot E[P_i|Q_i = 1]. \quad (4)$$

We also use conditional probability to derive an upper bound for the Quantity Upgrade Effect. The bound attributes the observed change in revenue entirely to the quality upgrade effect by setting the price paid by marginal consumers to zero. The formal derivation of the bound is presented in Appendix B.

4 Randomization Check

The experiment included several million users who visited the site over ten days. To check randomization, we test whether we can reject a 50% treatment assignment probability. Results are shown in Table I. While the odds of assignment to the treatment group are 50.11% in the full sample, the large scale of the experiment allows us to reject the null hypothesis of a 50% assignment probability at the 5% level. Upon closer scrutiny, we discovered two glitches in the randomization: first, all users who logged in during the first 30 minutes of the experiment were assigned to the treatment group. Second, users on a particular browser-operating system combination were also skewed to the treatment group. After eliminating these two groups we can no longer reject a 50% assignment at the 1% level.⁹ We therefore exclude these users in our main analysis.¹⁰ Although the probability of treatment remains slightly above 50%, the difference is economically insignificant.

Table I: Treatment Assignment

Sample	% Unidentified	% Site in Sample	% Back-end Fees	T-statistic
Full	0.78%	100%	50.11%	4.28
Time Restriction	0.78%	99.82%	50.09%	3.41
Time & Browser Restriction	0.82%	66.12%	50.06%	1.99

Notes: This table reports the assignment of StubHub users (cookies) to different treatment cells. Each row corresponds to a different sample restriction. The T-statistics are from a two-sided test with a null of a 50% assignment probability.

As a robustness check on randomization, we test whether UF and BF users share similar observable characteristics. Unfortunately, as treatment was assigned before users are required to log-in, the set of observables is limited. For example, we observe a user’s purchase history only if they log on to the site during the experiment or if they have not cleared their cookies

⁸Expected revenue using conditional probability is $E[R_i] = E[P_i|Q_i = 1] \cdot Pr\{Q_i = 1\} = E[P_i|Q_i = 1] \cdot E[Q_i]$.

⁹Or at the 5% level in a one-sided test against the null that the treatment assignment is $> 50\%$.

¹⁰However, our main results are robust to their inclusion in the sample.

after a recent visit. However, we do see site visits since the last cookie-reset, which we use to measure experience. We use this proxy as a left-hand side variable in specification (1). Row 1 of Table II shows that the two groups have almost identical experience levels. BF and UF users also visit the site at similar hours-of-the-day, and are equally likely to use mac computers (rows 2 and 3). These results give us confidence that the randomization was successful.

Table II: Covariate Balance

User Characteristic	% Difference	T-statistic
Experience	0.01	0.02
Hour	-0.08	-1.6
Mac User	0.16	0.01

Notes: This table presents summary statistics for differences between the BF (treatment) and UF (control) groups in our experiment from August 19 - 31, 2015.

5 Results

Our framework indicates that obfuscation should encourage consumers with a low willingness-to-pay for quality to switch from the outside option to purchasing a ticket on StubHub, and also encourage consumers to switch from purchasing lower to higher quality tickets. Table III column 1 shows the net effect on revenue of the price salience treatment. Consumers identified with cookies in the Back-end Fee group, where fees are obfuscated, spend almost 21% more than those assigned to the Upfront Fee group. We show revenue effects for the session (same-day) and over the entire experiment (10 days), and point estimates are large and statistically significant at the 1% level for both.

Unfortunately, quantifying salience is difficult, so it is hard to benchmark our estimate to Chetty et al. (2009). (While the change in user experience in the StubHub experiment is similar in spirit to their experiment of adding taxes to supermarket shelf prices, it is not clear how closely they align.) They find that obfuscating a 7.35% tax leads to an 8% revenue increase. On StubHub, obfuscating a 15% fee leads to a 21% revenue boost.¹¹ Our findings, detailed below, suggest that upgrades augment the salience effect in our setting.

5.1 Quantity Effect

We first examine the effect of salience on quantity. The third row of Table III shows that price obfuscation increased the transaction rate over the full course of the experiment

¹¹Fee documented in: Katy Osborn. September 1, 2015. "Why StubHub is tacking on ticket fees again." *The Wall Street Journal*.

by 14.1%. The second from last row shows within a cookie-session, consumers in the BF group are 12.43% more likely to purchase a ticket during a visit (the estimate is significant at the 1% level). Fees average roughly 15% of ticket prices, suggesting a per-session salience elasticity of $0.1243/0.15 = 0.87$, which is a similar order of magnitude to the elasticity of 1.1 found in Chetty et al. (2009). The 10-day elasticity is larger than the session elasticity ($0.141/0.15 = 0.94$), suggesting that the long-run effects of salience may be even greater.

Table III: Effect of Salience on Purchasing

	Back-end vs Upfront Fees % Difference	
	Baseline	Conditional on Purchasing
Cookie 10-day Revenue	20.64% (1.38)	5.42% (1.37)
Average Seat Price	—	5.73% (1.5)
Propensity to Purchase at Least Once	14.1% (0.09)	—
# Transactions within 10 Days	13.24% (0.88)	-0.9% (0.58)
# Seats within 10 Days	11.37% (1.17)	-2.32% (0.84)
12-Month Churn	—	-3.29% (0.66)
Cookie Session Revenue	18.96% (1.27)	5.61% (1.27)
Cookie Session Propensity to Purchase	12.43% (0.6)	—

Notes: This table presents estimates of how fee salience affects purchasing. Effects are presented as percent differences between treatment (BF) and control (UF) users, as per equation 1. Heteroskedasticity-robust standard errors are reported in parentheses. The sample in column 1 is all visitors to StubHub between August 19 and August 31 of 2015. Column 2 restricts to users who made at least one purchase during the same period.

Table III also provides estimates of how salience impacts the number of tickets purchased. Our framework ignores the consumer’s decision of how many seats to buy and describes a world where consumers need a fixed number of seats and either buy that exact number or buy none at all. In reality, of course, consumers might enlarge their parties if they perceive prices to be lower. To the contrary, we find that Back-end Fee users buy 2.4% fewer seats, conditional on making at least one purchase at StubHub. Admittedly, this effect is swamped by the increased probability of buying at least one ticket on StubHub, but hints at the nuance in salience responses. The lower number of seats suggests that the marginal consumers lured

by the Back-end Fee treatment buy slightly fewer tickets.¹²

5.2 Quality Upgrade Effect

The second column of Table III compares differences in the Back-end and Upfront Fee groups' behavior conditional on a purchase. This comparison allows us to assess how salience affects average purchase prices: BF users spend 5.42% more than their UF counterparts. From the platform's perspective, the combination of the Quantity Effect and the Quality Upgrade Effect implies that the effect of salience on their bottom line is substantially larger than suggested in the earlier literature, which did not consider product quality upgrades.

Using equation (4) we can calculate the increased revenues that are due separately to the Quantity effect and the Quality Upgrade Effect. From Table III, we observe that $\Delta P = 5.42P$ and $\Delta Q = 14.1Q$ and hence, rewriting Equation (4) without the expectations operator and subscripts for brevity,

$$\Delta R = \Delta P \cdot Q + \Delta Q \cdot P = 5.42 \cdot QP + 14.1 \cdot QP. \quad (5)$$

Dividing both the left- and right-hand sides of (5) by revenues, $R = QP$, we calculate the percent change in revenues ($\Delta R/R$) to be 19.52%, of which 5.42% (about 28% of increased revenues) are from the Quality Upgrade Effect. Note that the number of seats declines slightly, so that the change in the average purchase price per seat is even greater (5.73%).


We interpret the change in purchase price as evidence of an upgrade effect, where obfuscating fees leads consumers to buy more expensive, higher quality tickets. This finding is consistent with Lynch and Ariely (2000), who find that subjects in a lab experiment bought higher quality wine when prices were not displayed alongside product descriptions (and were only shown at checkout). Our framework indicates that the change in the average purchase price constitutes a lower bound for the upgrade effect – and while smaller than the quantity effect, even this lower bound is economically meaningful. Our upper bound calculation in (8) is 20.28%, suggesting that the Quality Upgrade Effect may even exceed the Quantity Effect.

We provide auxiliary evidence on the upgrade effect using data on seat locations. In particular, we examine whether Back-end Fee users bought seats closer to the stage. Rows are often labeled using letters, where letters earlier in the alphabet correspond to a better view.¹³ Conditional on purchasing a ticket, we separately calculate the probability that a BF and UF user purchases a seat in each row. Figure IV graphs the relative probability (the ratio of the two probability mass functions), along with 95% confidence intervals, which are calculated point-wise. Back-end Fee users are relatively more likely to purchase seats in rows A through D, which are the very first rows, and the likelihood declines for rows later in the

¹²A second possibility is that the revelation of fees at checkout induces Back-end Fee users to reduce the number of seats that they intend to purchase once they observe the fee-inclusive price.

¹³As numbering schemes vary across venues, letter position only proxies for quality.

Figure IV: Difference in Likelihood of Purchase by Row (BF versus UF users)



Notes: This figure plots the relative purchase likelihood by ticket row letter for users in the treatment (BF) and control (UF) groups. Letters earlier in the alphabet generally correspond to seats that are nearer to the event stage.

alphabet. These patterns provide further evidence of the Quality Upgrade Effect.

5.3 A Second Experiment: Event-Level Randomization

The 2015 experiment randomized salience across users so that BF and UF users had the same StubHub experience except for fee presentation—fees were included in the search results only for UF users. In an earlier experiment performed in 2012 at StubHub, fee salience was randomized at the event level, which presents distinct challenges, but offers a nice robustness check for the 2015 experiment.

First, StubHub’s unique inventory threatens the independence assumption for the 2015 experiment, but not for its 2012 counterpart. Suppose that price obfuscation merely accelerates, but does not actually alter, the consumer’s purchase decision. In this case, BF users will tend to buy early in the 2015 experiment, which may reduce inventory for UF users. Comparing purchase probabilities without taking this censorship into account would mistakenly indicate a positive treatment effect. In other words, treating user A affects user B (see Blake and Coey (2014) for a discussion of this challenge on eBay). Fortunately, the 2012 experiment does not suffer from the same contamination concern because all tickets for a particular event share the same treatment status.

A second challenge that the 2012 experiment addresses is multi-device use. In the 2015 experiment, we sort users into BF or UF the first time they touch an event page on StubHub

Table IV: 2012 Experiment Results
Back-end vs Upfront Fees

	% Difference
Purchase Probability	-12.38% (6.63)
Percentile of Choice Set Selected	-11.97% (5.62)

Notes: This table presents estimates of how fee salience affects customer purchasing based on data from the 2012 StubHub experiment, where salience is randomized at the event level. Effects are presented as percent differences between Back-end and Upfront Fee users. Standard errors are clustered at the event level and reported in parentheses.

during the experiment period. StubHub employs cookies to track users, so that the user remains in the appropriate group throughout the trial. However, cookies differ across devices, and a user would be re-randomized into the BF or UF group if she used a different device. Switching devices is particularly problematic if its incidence depends on initial treatment assignment. As an example, if UF users – upon seeing higher initial prices – delay their purchases and revisit StubHub on a second device, then the BF treatment would be positively correlated with purchasing. In the 2012 experiment, tickets to each event retain their treatment status regardless of the device that consumers use.


Finally, randomization at the event level provides insight into general equilibrium effects examined in section 7. We have shown that when StubHub alters the consumer’s experience, it alters sellers’ behavior. Salience might also affect price levels, which is hard to gauge given the unique inventory on StubHub. For example, if price obfuscation attracts more elastic buyers, then sellers might lower their prices. If these effects are large, then the 2015 experiment does not provide the true counterfactual of interest: what happens when all users face BF? Instead, the econometrician only observes what happens on StubHub when fees are shrouded for 50% of users. The 2012 experiment answers this question because a ticket-seller for a particular match faces an entirely BF or UF audience, but not a mix of both.

In the 2012 experiment, 33 out of 99 Major League Soccer Games were randomly selected for UF. Prices for tickets to these games included fees, even from the initial event page. The remaining 66 matches had the BF experience, which at the time was the site-wide user experience. The results from the 2012 experiment, displayed in table IV, confirm our 2015 findings: fee salience reduces revenue substantially. Consumers are 13% less likely to buy tickets to an Upfront Fee match.¹⁴ The difference has a p-value of 0.076, with standard errors clustered at the event-level.

We also examine whether users upgrade to more expensive tickets for BF games. Un-

¹⁴ Note that fees were approximately 10% in 2012.

Figure V: Percentile of Choice Set Purchased in the 2012 Experiment



Notes: This figure plots the PDF of purchases by price percentile separately for treatment (BF) and control (UF) users on StubHub.com. To calculate price percentiles, we reconstruct the set of available tickets on StubHub.com during each user’s site visit.

fortunately, tests based on purchase prices are under-powered because of the high sampling variance across matches. To control for the unobserved popularity of matches, we test whether users purchase from the same quantile of price in BF versus UF matches. For each transaction, we calculate where the purchase ranks in a user’s choice set (StubHub’s entire inventory for the match at the time of purchase). On average, consumers buy from a 12% lower quantile for UF compared to BF games. Figure V shows the full distribution of purchase quantiles for BF and UF matches.

While these results are heartening, we prefer the 2015 experiment for its larger sample size. Further, experimentation at the event-level suffers from a different kind of contamination bias: consumers may substitute away from UF matches (which appear more expensive) to BF matches. The 2015 experiment is not vulnerable to this type of contamination. Another complementarity between the two experiments is that they differ in initial conditions: in early 2012, StubHub used a BF policy, while in 2015, the site used a UF policy. Our results suggest the effect of price salience at StubHub is similar despite the difference in the status quo. The ability to execute two experimental designs is one advantage of the StubHub setting.

Table V: Purchase Funnel Behavior by Fee Salience

	Percentage Click Through from Prior Page			Average Ticket Price		
	<u>BF</u>	<u>UF</u>	<u>% Difference</u>	<u>BF</u>	<u>UF</u>	<u>% Difference</u>
Event Page	–	–	–	\$1.00	\$0.84	18.73%
Ticket Details	27.96 (0.01)	23.56 (0.01)	18.67% (0.00)	\$0.86	\$0.78	10.16%
Review & Submit	–	–	–	\$0.56	\$0.52	7.44%
Purchase	18.52 (0.06)	33.41 (0.1)	-44.58% (0.00)	\$0.42	\$0.39	6.57%

Notes: This table reports means and standard errors (in parentheses) of user behavior in the StubHub purchase funnel. Average ticket prices are normalized by the average price of tickets selected by Back-end Fee users on the Event Page.

6 Mechanisms

6.1 Misinformation

In this section, we leverage StubHub’s detailed data to better understand why fee salience affects consumers so greatly. First, we examine consumer misinformation using web-browsing behavior. If consumers do not anticipate fees, then they will receive a negative surprise at check-out and should be more likely to exit when the fee first appears. For consumers who are nearly indifferent between purchasing at the base ticket price, the fee makes the outside option their utility-maximizing choice. Importantly, a misinformation theory offers implications about where (in the purchase funnel) Back-end and Upfront Fee users will differentially exit.

To buy a ticket, a user follows StubHub’s “purchase funnel” on the website as follows: (1) the consumer first sees the event page, which contains a seat map and a sidebar with top ticket results, sorted by price in ascending order; (2) once a consumer clicks on a ticket, the ticket details page appears; (3) the consumer proceeds to the checkout page where a final purchase decision is made; (4) the purchase confirmation page completes the process.¹⁵ BF users are shown lower prices than their UF peers until stage (3), when they are shown the final price, inclusive of fees. If consumers are ignorant of fees, there should be a larger drop off between stages (1) and (2) for the UF group, since they see higher prices initially. But there should be a larger drop-off between stages (3) and (4) for the BF group. If the former is larger than the latter, then Back-end Fees increase the quantity sold.

The left panel of Table V shows the absolute and relative rate of UF and BF user arrivals between these key steps in the purchase process. Consistent with misinformation, Back-end

¹⁵Before reaching the Checkout page, a log-in page appears unless the consumer was already logged into their account. Many searches are non-linear, where consumers examine multiple event pages (see Blake et al. (2016)). BF users might even return to stage (1) once they see the additional fees leveled at stage (4).

Fee users are almost 19% more likely to select tickets (transition from stage 1 to 2) than Upfront Fee users. The difference is statistically significant at the 1% level and economically large. In contrast, the drop off rate at the final stage (purchase) is much larger for BF users, as they are almost 45% less likely to purchase at checkout.

The right panel of Table V presents the average selected ticket price at each step in the purchase funnel for a subset of events. The average price of tickets under consideration declines at each step, suggesting that quality also drops. As the theory predicts, UF users always select cheaper tickets than BF users, but the difference narrows as users move closer to purchase. When fees are revealed, the gap is just under 7% compared to an initial difference of almost 19%. In sum, BF users are more likely to contemplate buying expensive tickets but when fees are revealed, more of the (potentially surprised) BF users exit than the UF users who see no change in their expected outcome.

One important question, from both the firm’s and a policy maker’s perspective, is whether consumers learn about the fees over time. As an example, consumers could act as if they do not anticipate fees in their ticket selection each time they visit the site. In this case, websites stand to gain substantially by shrouding fees. This implication contrasts with a model where consumers anticipate a fee, but do not know the exact level. In a model with learning, once a consumer makes a purchase, she updates her priors on future StubHub fees and does not make the same ‘mistake’ twice.

To examine learning, we repeat our principal analysis (Table III) separately by level of user experience. If consumers learn, then experience ought to lessen the response to obfuscation. Of course, experience is endogenous, so experienced users may react differently to salience for other reasons (as an example, they may be higher income). Nonetheless, examining responses across experience groups hints at how learning might work in this setting.

To measure experience, we calculate the number of visits each cookie has made to StubHub prior to the experiment. A 2006 ComScore study found that 31% of users clear their cookies within 30 days, so we interpret this as a short-term measure of experience.¹⁶ Unfortunately, we cannot exploit information about logged-in users (like number of past transactions) because log-in is a potential response to our treatment; users who see lower prices initially may be more likely to log in to the website in order to purchase. Our measure does capture the most recent interactions with StubHub, which are likely to be the most relevant for a user’s knowledge of the site.

We hypothesize that frequent StubHub users ought to be aware of fees and therefore less sensitive to salience. We split users into three groups: new users (no recorded visits), low experience (1-9 visits), and high experience (10 or more visits). Table VI shows that the treatment effect is smaller for cookies with at least 10 site visits: the revenue effect is 15% compared to 21%. These results suggest that salience may be most important in markets

¹⁶<https://www.comscore.com/Insights/Blog/When-the-Cookie-Crumbles>

Table VI: Saliency by User Experience

	% Difference		
	<u>New User</u>	<u>Low Experience</u>	<u>High Experience</u>
User 10-day Revenue	21.52% (1.92)	21.80% (2.29)	15.09% (4.4)
Propensity to Purchase at Least Once	15.33% (0.653)	13.68% (1.15)	10.19% (2.42)
# Transactions within 10 Days	14.33% (1.17)	13.53% (1.23)	8.81% (2.94)
% Sample	67%	27%	6%

Notes: This table reports coefficient estimates of how fee saliency affects purchasing (equation 1) for users of different experience levels. Estimates are presented as percent differences between treatment (BF) and control (UF) users. Heteroskedasticity-robust standard errors in parentheses. See table III for pooled estimates.

where consumers purchase infrequently (for example, real estate or automobile markets). However, effects are still large for the most experienced group (the top 6% of users), which indicates only limited consumer learning. Because experience is not randomly assigned in the population, we interpret this evidence as suggestive, rather than causal.

We examine user churn to understand the long-run effects of saliency. If obfuscation preys on misinformation, then marginal BF consumers, who would not purchase if shown fees upfront, may be more likely to abandon StubHub after seeing fees for the first time. Unfortunately, we cannot identify marginal consumers among the pool of BF consumers. We also cannot compare the return rates of all BF and UF users, as there is no way to track future purchases of users who do not log-in to the site. Instead, we compare the return rates of BF and UF users who purchase during the experiment. As Table III shows, BF users are 3.3% less likely to churn, which is inconsistent with the simple misinformation story. We emphasize caution in interpreting churn, however, as it potentially confounds multiple treatments: BF users may learn about the platform fees when they make a purchase, but they may also learn about StubHub’s reliability, speed, quality, etc. This additional learning may increase a consumer’s likelihood of purchase, even if obfuscation effects are short lived.

As a robustness check, we compare the likelihood of return for consumers who were logged into StubHub *before* the experiment. We can track these users’ purchases after the experiment’s conclusion, regardless of whether they made a purchase during the experiment window. The difference between BF and UF return rates drops to 0.65% and loses statistical significance. While this sample contains consumers with high attachment to StubHub, this comparison also indicates that saliency effects persist beyond initial misinformation.

Finally, to shed light on the persistence of saliency, we construct a panel dataset that tracks the purchases of BF and UF users over a six month period centered around the

experiment window (May 18 through December 1, 2015). We have already established that BF users spend more, conditional on purchasing, during the experiment. On September 1, the entire site switched to BF, so that the only difference between users who *had been* assigned to BF versus UF is their experience with the back-end fees. If salience effects are short-lived, then we would expect UF users, who now experience back-end fees for the first time, to outspend their BF counterparts who have 10 days of experience. On other hand, if salience effects persist, then the UF-BF difference should dissipate after the experiment, as both groups spend more than they would have in an UF environment. If i denotes the user and t the purchase date, we model purchase price using the following specification:

$$\ln p_{it} = \alpha_0 + \sum_{w=1}^W \alpha_w \cdot 1\{\text{week}_t = w\} + \sum_{w=1}^W \beta_w \cdot 1\{\text{week}_t = w\} \times T_i + \epsilon_{it} \quad (6)$$


where $1\{\text{week}_t = w\}$ is an indicator that the purchase occurred during week t in our sample and T_i is a treatment indicator. For ease of interpretation, the week 14 indicator is labeled *experiment* and comprises 10 rather than 7 days. Purchases the first day of and after the experiment are omitted to account for any engineering lags in the user interface switch. We estimate equation (6) using the sample of users who purchase during the experiment window because these are the only users we can reliably track. During the purchase process, users log into the site, allowing us to identify their prior and subsequent purchases. Standard errors are clustered at the user level to account for serial correlation in individual purchasing decisions.

Figure VI displays the estimates of the interactions between the BF treatment indicator and each time period. BF and UF users spend similar amounts before the experiment, when both groups experience UF. As in Table III, we find that during the experiment, BF users spend almost 6% more than UF users, conditional on purchasing at least one ticket. However, in the three month period following the experiment, when all users experience BF, there is no difference in spending between the two groups. The results are robust to the inclusion of both buyer and day fixed effects. These event study findings, taken together with our results on experienced users and churn, indicate that salience effects are persistent. They suggest that users do not learn to anticipate the correct fee level after going through the purchase funnel with back-end fees at least once.

6.2 Consideration Sets and Search Frictions

In this section, we present evidence on forces beyond misinformation that might contribute to the importance of salience: consideration sets and search frictions. First, we consider whether fee obfuscation widens users' consideration sets. A growing body of literature (e.g. Goeree (2008)) suggests that potential consumers often ignore a large fraction of inventory,

Figure VI: Spending Before, During, and After the Platform Switch to Back-end Fees




Notes: This table includes data on purchases between 05/18/2015-12/01/2015, excluding 08/19/2015 (the first day of the experiment) and 09/01/2015 (the first day after the experiment). The data includes only those customers who purchased *at least* once during the experiment window.

and instead focus on choosing between a few products. StubHub presents inventory to consumers in ascending price order, so that expensive tickets are not visible to the consumer unless she actively scrolls down or filters the results (e.g., by section). It is possible that obfuscating fees might draw user attentions to a wider array of products, leading BF users to make different purchase decisions than their UF counterparts. We find that BF users scroll 10% more often, a difference that is statistically significant at the 1% level.

When fees are revealed, BF consumers are already at check-out with their tickets, but they may go back to the event page to re-optimize and purchase cheaper seats. We find that less than a quarter of BF users exercise this option, which is consistent with a search friction beyond misinformation. Figure VII shows the average number of tickets viewed by BF and UF users. BF cookies are 56% more likely to view multiple ticket listing compared to their UF counterparts. Table VII shows that BF users view cheaper tickets upon their return to the listings page from the checkout page (six percentage points cheaper). In contrast, UF users, who are less likely to return overall, view more expensive tickets if they do.

Figure VII shows that BF users are twice as likely to view three or more listings than their UF counterparts. Viewing more than two tickets suggests the effects of price obfuscation extend beyond an initial confusion about fees. BF consumers who return to the event page have seen fees for their initial selection, but they must calculate the StubHub fee for each new ticket they consider. If calculation costs are high, as hypothesized by Morowitz et al. (1998) or Ellison and Ellison (2009), consumers might choose to go down the funnel multiple times rather than compute the fees themselves. Obfuscation as a search friction is consistent

Figure VII: Number of Listings Viewed by Fee Salience



Notes: This histogram plots the number of listings viewed across users. The distribution is plotted separately for treatment (BF) and control (UF).

with our findings on experienced customers, who ought to anticipate fees but might still bear a higher search cost when fees are hidden. This evidence is in line with Ellison and Ellison (2009), who find that firms endogenously create such frictions to soften price competition.

7 Two-Sided Responses


In this section, we provide evidence on the effect of fee salience beyond changes in consumer behavior. Note first that in two-sided markets like ticket resale, changes to the buyer experience may spill over onto sellers. As an example, if obfuscation lifts seller profits (by increasing buyer spending), then more sellers may enter the marketplace. In turn, increased seller participation may bolster competition and help buyers. These sorts of externalities complicate welfare analyses in two-sided markets.

Table VII: Average Price of Tickets Viewed Relative to UF Initial Selections

Back-end Fees		Upfront Fees	
Initial Checkout	Follow-up Actions	Initial Checkout	Follow-up Actions
8.3%	0.8%	0.0%	1.8%
(1.9)	(1.2)	(-)	(0.6)

Notes: This table reports means and standard errors for the relative price of tickets viewed across the treatment and control groups. Estimates are normalized by the price of tickets initially brought to check-out by UF users.

Figure VIII: Fraction of Listings by Row Letter



Notes: This figure plots the number of listings by row letter relative to base rows U-Z. The two vertical red lines denote the start and end of the 2015 fee salience experiment.

7.1 Ticket Quality

As a first step, we examine whether inventory responds to the use of BF pricing, with a focus on ticket quality. Section 5.2 shows that buyers upgrade to higher quality seats when fees are less salient, making StubHub a more attractive platform to sellers of high quality tickets. Figure VIII shows the evolution of inventory on StubHub over time by row letter. Visual inspection suggests that the relative number of seats in front rows (A-E) compared to back rows (U-Z) increase after the switch to BF. Consistent with Ellison and Ellison (2009), we find that sellers respond to the change in the buyer experience.

To further investigate seller responses, we test for a break in listing quality during and after the experiment, when the whole site switched to BF. To measure quality, we construct a row number variable, *Position*, which counts the number of rows in between the seat and row A plus one (taking a value of one for seats in row A). We then construct an event study, where the log number of listings is the dependent variable. We are interested in the coefficient on the interaction between $\ln(\text{Position})$ and an indicator for the post period as below:

$$\ln \text{Listings}_{it} = \beta_0 + \beta_1 \cdot \ln(\text{Position}_{it}) + \beta_2 \cdot \text{Post}_t \times \ln(\text{Position}_{it}) + \Gamma_t + \epsilon_{it} \quad (7)$$

Our preferred specification includes day fixed effects, Γ_t , which control for any site-wide fluctuations that affect all types of tickets simultaneously. Columns 1 & 2 in table VIII present the coefficient estimates on the interaction terms, which are negative and statistically significant at conventional levels for both the experiment and its aftermath. The point estimates imply that a ticket listed on StubHub is 3.7% more likely to be in row A than row B following the experiment (under BF) compared to before (under UF). The increase in high quality listings underscores the complexity of platform design, as changes to one side of the market influence the entry decisions on the other.

Table VIII: Changes in Listings Following Back-end Fees

	Log Number of Listings			
	(1)	(2)	(3)	(4)
Log Position \times Post	-0.123 (0.020)	-0.123 (0.017)		
Round \times Post			0.315 (0.064)	0.315 (0.018)
Date FE	No	Yes	No	Yes
Observations	4,680	4,680	360	360

Notes: Heteroskedasticity-robust standard errors in parentheses. Data from 06/01/2015 - 12/01/2015 at the daily level. Controls include log Position (the letter’s position in the alphabet, where A occupies the first position) in columns 1 & 2 and an indicator for a round base price in columns 3 & 4. Column 1 & 3 also include an indicator for the post period.

7.2 Ticket Prices


Second, we consider whether prices respond to Back-end Fees. Ideally, we could test whether Back-end Fees induce sellers to increase or decrease prices by comparing price levels before and after the site switches from UF to BF in September 2015. However, this time series variation is confounded by changes in site inventory over time. The challenge is that the tickets listed and sold in August differ from those listed and sold in September because different events are held in the two months. As an example, the 2015 NFL season kicked off on September 10th. Instead of examining price levels, we focus on another aspect of pricing: the use of round numbers.

An extensive literature in marketing documents the appeal of round number pricing (amounts that end in zeros or nines).¹⁷ If sellers aim to employ round number pricing, then they ought to adjust prices in response to the site’s switch from UF to BF. That is, under UF, a seller should set their list (or “base”) price so that the fee-inclusive price (list price + buyer fee) that is shown to the consumer is round. In contrast, under BF, the seller should

¹⁷For example Monroe (1973) or more recently Backus et al. (2019).

set a round list price. Thus, we examine whether sellers are more likely to set base prices at round numbers after the switch to Back-end Fees.

Figure IX: Percent of Listings with Round Prices



Notes: This figure plots the fraction of StubHub listings with round base prices for a six month window around the 2015 fee salience experiment. The two vertical red lines denote the start and end date of the experiment. The sample comprises listings that were created or modified each day.

As shown in figure IX, the share of listed tickets with round base prices increases by approximately 5 percentage points following the switch to Back-end Fees. To be transparent, we examine only the prices of listings that were added or modified on each date, and we categorize prices that end in “.00” or “.99” as round. Columns 3 & 4 in table VIII present results of the regression analogue of figure IX, where we adopt specification 7 so that the independent variable of interest is an indicator for a round listing price. The results indicate an economically and statistically significant increase in the use of round listing prices following the switch to BF. This trend shows that sellers adjust their pricing policies in response the buyer’s experience, which is consistent with Ellison and Ellison (2009).

8 Discussion

As the online share of transactions continues to grow, so too does the scope for regulations that guarantee the efficient functioning of markets. Chief among proposed regulations has been increasing the transparency of mandatory fees. Using data from a randomized control trial on StubHub, we find that shrouding buyer fees increases total revenue by about 20%. In

the experiment, the control group was shown fee-inclusive prices from the initial search page, while the treatment group was shown base prices until the checkout page. We decompose the impact of obfuscation into a quantity effect and a quality effect. The latter accounts for at least 28% of the revenue bump because consumers upgrade to higher quality products when they observe lower prices initially. We find that consumers who are shown fees upfront drop off early in the purchase funnel, while those shown fees later are more likely to exit after the site displays total prices, consistent with consumer misinformation.

We find that salience persists beyond initial misinformation. Experienced users, who arguably should anticipate the fee, spend 15% more on StubHub when the fee is shrouded. More strikingly, after the platform switched to Back-end Fees, the users exposed to the BF treatment during the experiment spend similar amounts to those newly exposed to Back-end Fees. This behavior suggests that short-term experience with Back-end Fees does not give users an advantage in anticipating true final prices. These patterns indicate that salience is not a one-off phenomenon, which becomes irrelevant as consumers learn about the sales environment. It is perhaps unsurprising, if not reassuring, that we find that sellers respond to changes in the salience of the buyer experience. Sellers are more likely to list high quality tickets and to use round number prices when fees are presented at the back-end, highlighting the nuance of salience effects on a platform.

Our results also demonstrate that price salience looms large in markets where consumers purchase only intermittently. The existing literature focuses on contexts where consumers purchase frequently, such as grocery stores in Chetty et al. (2009). In these settings, consumers plausibly hold strong beliefs about both the amount and presentation of fees and taxes,¹⁸ and so we might interpret their response to an abrupt change in salience as a reaction to off-equilibrium path play. In contrast, most users who visit StubHub during our experiment are new to the site. Their reactions to salience may more closely parallel reactions in markets like real estate, higher education, or automobiles, where policymakers may wish to mandate fee disclosure.¹⁹

¹⁸Chetty et al. (2009) provide survey evidence that the modal consumer in their setting identifies the correct tax level.

¹⁹E.g., starting in 2012, the Department of Transportation required airlines to advertise fee-inclusive prices.

References

- Backus, Matthew, Thomas Blake, and Steven Tadelis**, “On the empirical content of cheap-talk signaling: An application to bargaining,” *Journal of Political Economy*, 2019, *127* (4), 1599–1628.
- Blake, Thomas and Dominic Coey**, “Why Marketplace Experimentation is Harder than it Seems: The Role of Test-Control Interference,” *Proceedings of the 2014 ACM Conference on Economics and Computation*, 2014, *EC15*, 567–582.
- , **Chris Nosko, and Steven Tadelis**, “Returns to Consumer Search: Evidence from eBay,” *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016, *EC17*, 531–545.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Salience and Consumer Choice,” *Journal of Political Economy*, 2013, *121* (5), 803–843.
- Brown, Jennifer, Tanjim Hossain, and John Morgan**, “Shrouded Attributes and Information Suppression: Evidence from the Field,” *Quarterly Journal of Economics*, 2010, *125* (2), 859–876.
- Brynjolfsson, Erik and Michael D. Smith**, “Consumer Decision-making at an Internet Shopbot: Brand Still Matters,” *Journal of Industrial Economics*, 2001, *XLIX* (4), 541–558.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *The American Economic Review*, September 2009, *99* (4), 1145–1177.
- DellaVigna, Stefano**, “Psychology and Economics: Evidence from the Field,” *Journals of Economics Literature*, 2009, *47* (2), 315–372.
- Dertwinkel-Kalt, Markus, Mats Köster, and Matthias Sutter**, “To Buy or not to Buy? Shrouding and Partitioning of Prices in an Online Shopping Field Experiment,” *CESifo Working Paper*, 2019.
- Ellison, Glenn**, “A Model of Add-On Pricing,” *The Quarterly Journal of Economics*, May 2005, *120* (2), 585–637.
- and **Sara Fisher Ellison**, “Search, Obfuscation, and Price Elasticities On The Internet,” *Econometrica*, March 2009, *77* (2), 427–452.
- Feldman, Naomi E. and Bradley J. Ruffle**, “The Impact of Including, Adding, and Subtracting a Tax on Demand,” *American Economic Journal: Economic Policy*, 2015, *7* (1), 95–118.
- Finkelstein, Amy**, “E-Z Tax: Tax Salience and Tax Rates,” *The Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.
- Gabaix, Xavier and David Laibson**, “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, May 2006, pp. 505–539.
- Goeree, Michelle Sovinsky**, “Limited Information and Advertising in the U.S. Personal Computer Industry,” *Econometrica*, September 2008, *76* (5), 1017–1074.
- Greenleaf, Eric A., Eric J. Johnson, Vicki G. Morowitz, and Edith Shalev**, “The Price Does Not Include Additional Taxes, Fees, and Surcharges: A Review of the Research on Partitioned Pricing,” *Journal of Consumer Psychology*, 2016, *26* (1), 105–124.

- Heidhues, Paul and Botond Köszegi**, “Behavioral Industrial Organization,” in B. Bernheim, S. DellaVigna, and D. Laibson, eds., *Handbook of Behavioral Economics: Applications and Foundations 1*, Vol. 1 2018, chapter 6, pp. 517–612.
- Hossain, Tanjim and John Morgan**, “...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay,” *Advances in Economic Analysis & Policy*, 2006, 6 (2), 1–30.
- Lynch, John G. Jr. and Dan Ariely**, “Wine Online: Search Costs Affect Competition on Price, Quality, Distribution,” *Marketing Science*, 2000, 19 (1), 83–103.
- Monroe, Kent B.**, “Buyers’ Subjective Perceptions of Price,” *Journal of Marketing Research*, 1973, 10 (1), 70–80.
- Morowitz, Vicki G., Eric A. Greenleaf, and Eric T. Johnson**, “Divide and Prosper: Consumers’ Reactions to Partitioned Prices,” *Journal of Marketing Research*, November 1998, 35 (4), 453–463.
- Seim, Katja, Maria Ana Vitorino, and David M. Muir**, “Drip Pricing When Consumers Have Limited Foresight: Evidence from Driving School Fees,” *Working Paper*, February 2017.
- Taubinsky, Dmitry and Alex Rees-Jones**, “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment,” *Review of Economic Studies*, 2018, 85 (4), 2462–2496.

A A Model of Consumer Choice with Limited Fee Salience

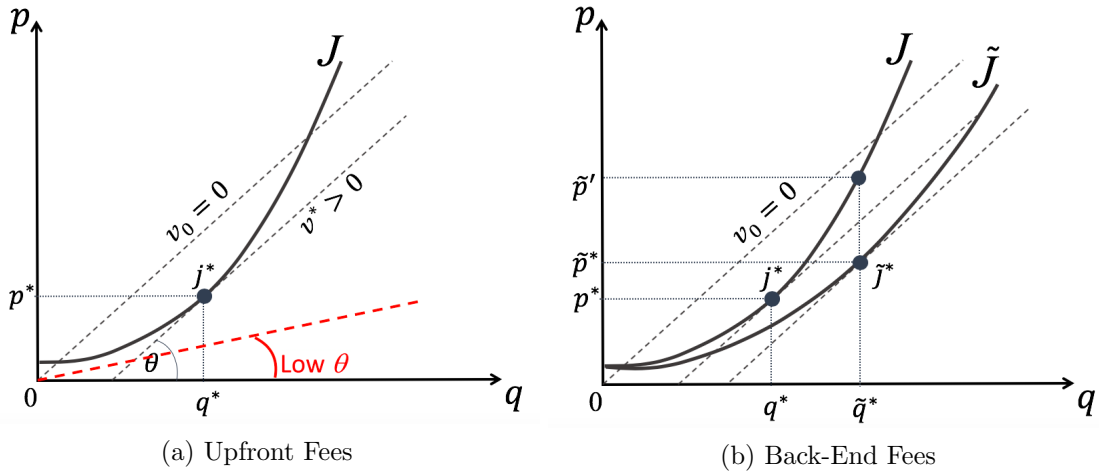
Consider a consumer who makes purchase decisions under two regimes. In the first, which we call Upfront Fees (UF), the final purchase price including all fees is shown to consumers when they browse the set of available tickets. In the second, which we call Back-End Fees (BF), consumers observe only list prices when they browse available products, and fees are revealed only after a particular ticket is selected for purchase.

First, we consider a consumer's choice when she observes fees upfront. She is presented with a convex and compact set of available tickets J , where her utility v_j from ticket $j \in J$ depends on its price p_j and quality q_j (e.g. section and row, delivery method, etc.) as follows:

$$v_j = \theta q_j - p_j.$$

The consumer's willingness to trade off quality for money is captured by her type $\theta \in [0, \bar{\theta}]$. Let 0 denote the outside option, with $q_0 = p_0 = 0$. Figure Xa illustrates her optimization problem: the set J of available tickets lies on and above the curved line, and the dashed line $v_0 = 0$ marks the consumer's indifference curve from not purchasing. The consumer chooses the ticket $j^* \in J$ on her highest indifference curve, yielding a utility of $v^* > 0$. A higher θ consumer will purchase a higher quality ticket at a higher price. For consumers with low enough values of θ (less steep indifference curves in Figure Xa), their indifference curve $v_0 = 0$ lies fully below the set J , and they will not purchase any ticket. It therefore follows that given a set of tickets J , there exists a threshold type $\underline{\theta} > 0$ such that a consumer of type θ will purchase a ticket if and only if $\theta > \underline{\theta}$.

Figure X: Optimal Ticket Choice



We model consumer optimization with Back-end Fees as a shift in the boundary of J . Namely, her choice now depends on the *perceived* price \tilde{p}_j of ticket j rather than its true final price. This is akin to reducing the salience of prices relative to quality as in Bordalo et al.

(2013) and is also similar to the way Finkelstein (2009) models salience. The consumer then selects $j \in J$ to solve her optimization problem:

$$\max_{j \in J} \tilde{v}_j = \max_{j \in J} \theta q_j - \tilde{p}_j$$

where the perceived price of not purchasing a ticket is also zero, $\tilde{p}_0 = p_0 = 0$. The established view on price salience is that $\tilde{p}_j < p_j$. That is, when fees are obfuscated, prices appear lower to consumers than they actually are, as illustrated in Figure Xb. The true price-quantity frontier is still J , however when the consumer chooses a ticket for purchase she perceives the frontier to be \tilde{J} , choosing the ticket \tilde{j}^* which has quality \tilde{q}^* and *perceived* price \tilde{p}^* .

Upon reaching the checkout and purchase phase, the ticket’s *actual* price—including all fees—is revealed to be $p' > \tilde{p}^*$. We assume, however, that the consumer will continue with the purchase at this final stage of the purchase funnel rather than go back to the selection stage with a newfound understanding that the true choice set is J .²⁰

Recall that the set of consumers with $\theta < \underline{\theta}$ will prefer not to purchase if they perceive the set of tickets to be J . Some of these consumers, however, will select a ticket for purchase if they perceive the set of tickets to be \tilde{J} . It follows immediately that there exists a threshold type $\tilde{\theta} \in [0, \underline{\theta}]$ such that a consumer of type θ will purchase a ticket if and only if $\theta > \tilde{\theta}$. Hence, the analysis above implies that fee obfuscation has two effects on consumer choice:

(1) Quantity Effect: *Under the BF treatment, a consumer is more likely to purchase.*

This prediction is consistent with the existing literature: more salient fees reduce the likelihood of purchase. However, it precludes at least two alternative effects of salience: first, if consumers anticipate fees (or hold unbiased beliefs) then perceived prices may not be lower than actual prices. Second, it is also possible that price obfuscation generates a “disgust” factor, wherein last-minute fees upset consumers. In that case, the quantity effect could be negative, contravening the standard price salience model.

When true final prices are higher than perceived prices and the difference is *increasing* in the listing price, then the model generates a second prediction: customers buy higher quality items than they would under the Upfront Fee regime. This condition would be satisfied, for example, if consumers simply ignored or underestimated a proportional fee or tax. More formally, for any ticket j , let \tilde{p}_j be the perceived BF price excluding fees and let p'_j be the true final price observed at checkout. We have,

(2) Quality Upgrade Effect: *If $p'_j - \tilde{p}_j > 0$ and $p'_j - \tilde{p}_j$ is increasing in q_j , then consumers buy higher quality tickets under BF.*

²⁰Several frictions could prevent consumers who reach checkout from going back to purchase a different ticket, such as loss aversion or the anticipation of re-optimization costs (e.g. having to calculate the fee for each set of tickets). We remain agnostic as to which of these best explain why consumers do not re-optimize, which is what we find in the data.

Conditional on purchasing, consumers upgrade to higher quality tickets under Back-end Fees and therefore spend more on the site. The earlier salience literature overlooks this effect, perhaps because previously studied settings offered little vertical product differentiation (e.g. Electronic Toll Collection systems as in Finkelstein (2009) or supermarket beauty aids as in Chetty et al. (2009)). Indeed, the log-log demand specification favored by earlier work leaves no scope for quality upgrades.

The Quality Upgrade Effect emphasizes how identification strategies must respect the impact of salience on quality choice. Consider the alcohol sales analysis of Chetty et al. (2009). They compare an excise (lump sum) tax to a sales (percentage) tax. The excise tax should arguably not effect the quality of beer chosen (conditional on purchase), since it makes each can of beer “in the choice set” more expensive by the same amount. The sales tax, however, may affect both the quantity and quality margins, since it is a percentage of the price. Simple comparisons of the revenue effects of excise and sales tax salience may therefore lead to inconclusive results.

B An Upper Bound for the Quality Upgrade Effect

We derive an upper bound for the Quality Upgrade Effect by setting the purchase price among marginal consumers to zero. That is, we assume that users who buy under BF but abstain under UF get tickets for free under the BF treatment. Formally, consider the following expression for the expected purchase price under Back-end Fees:

$$\begin{aligned}
E[P_i|Q_{i1} = 1, T_i = 1] &= E[P_i|Q_{i0} = 1, Q_{i1} = 1, T_i = 1] \cdot \frac{P\{Q_{i0} = 1\}}{P\{Q_{i1} = 1\}} \\
&\quad + E[P_i|Q_{i0} = 0, Q_{i1} = 1, T_i = 1] \cdot \left(1 - \frac{P\{Q_{i0} = 1\}}{P\{Q_{i1} = 1\}}\right) \\
&= (QUE + E[P_i|Q_{i0} = 1, T_i = 0]) \cdot \frac{P\{Q_{i0} = 1\}}{P\{Q_{i1} = 1\}} \\
&\quad + \underbrace{E[P_i|Q_{i0} = 0, Q_{i1} = 1, T_i = 1] \cdot \left(1 - \frac{P\{Q_{i0} = 1\}}{P\{Q_{i1} = 1\}}\right)}_{>0}
\end{aligned}$$

The first equality follows from a conditional probability decomposition of $E[P_i|Q_{i1} = 1, T_i = 1]$. Note that it also relies on choice monotonicity, which implies that $Pr\{Q_{i0} = 1|Q_{i1} = 1\} = \frac{Pr\{Q_{i0}=1\}}{Pr\{Q_{i1}=1\}}$. In the second equality, we add and subtract an additional term to create a term including QUE. This last equality contains two expressions, the second of which includes the expected price of tickets bought by the marginal users who buy under BF but abstain under UF,²¹ which we cannot observe but is greater than zero. If we assume that these consumers

²¹The types $\theta \in [\tilde{\theta}, \theta]$ in the model we present in Appendix A.

buy at a price of zero, thereby setting this last term to zero, we obtain the following upper bound for QUE:

$$\text{QUE} \leq E[P|Q_{i1} = 1, T_i = 1] \cdot \frac{\Pr\{Q_{i1} = 1\}}{\Pr\{Q_{i0} = 1\}} - E[P_i|Q_{i0} = 1, T_i = 0] \quad (8)$$


Importantly, all of the terms on the right-hand-side in equation (8) can be estimated directly from the data.

C Competition with Other Platforms

An additional consideration is how fee presentation at StubHub affects the broader competitive environment, including prices and inventory on rival sites. We focus on Ticketmaster and SeatGeek, two alternative secondary market for tickets, with Ticketmaster serving as the primary market for certain sporting and music events. At the time of the 2015 experiment, both sites employed Back-end Fees. It is possible that in comparison, StubHub appeared more expensive to consumers (because its listing prices included fees) and therefore less attractive to sellers. Thus, when StubHub itself switched to Back-end Fees in September 2015, it may have drawn sellers and buyers who would otherwise have frequented a rival platform. Unfortunately, we do not have access to listing or sales data from Ticketmaster or SeatGeek, so we investigate the effect of StubHub’s switch to Back-end Fees using data from GoogleTrends on queries.

Figure XI shows the evolution of queries over three years from September 2014 to September 2016. (To be clear, Google normalizes weekly query volume separately for each platform by dividing by the site’s peak over 2012-2017, so that the index ranges from 0 to 100 for each site. Queries for Ticketmaster are virtually flat, indicating that there is no effect of StubHub’s switch to BF. During the entire period, SeatGeek seems to be gaining popularity, but again, there is no evidence of a trend break in September 2015 when StubHub makes the change. We formally test for a change in Ticketmaster and SeatGeek queries by adapting specification 7 so that the right-hand side interactions are with indicators for Ticketmaster and SeatGeek (rather than Position) and the left-hand side variable is the Google query index. The omitted category is queries for StubHub itself. Table VIII presents results, which shows an economically and statistically insignificant change in searches for Ticketmaster. In contrast, the coefficient on the interaction between SeatGeek and the post indicator is positive and statistically significant in columns 1 & 2, where the latter includes date fixed effects. To accommodate the gradual increase in SeatGeek queries during this period visible in figure XI, we add a site-specific time trend in column 3; the coefficient on the interaction term for SeatGeek and the post indicator halves in magnitude and reverses sign. Our interpretation of these results is that they provide little evidence that other ticket resale platforms were

Figure XI: Google Queries for Competing Ticket Resale Platforms



Notes: This figure plots the google trend index for StubHub, SeatGeek, and Ticketmaster for a two year window around the fee salience experiment. The index is normalized separately for each site based its peak over 2012-2017.

affected by StubHub's switch to Back-end Fees. More work with data that speak to rivals' sales and not simply queries is needed, however, to give a definitive answer.

Table IX: Changes in Google Searches Following Back-end Fees

	Google Queries Index		
	(1)	(2)	(3)
Ticketmaster \times Post	0.019 (2.691)	0.019 (2.408)	-1.092 (4.877)
SeatGeek \times Post	15.827 (2.691)	15.827 (2.155)	-8.765 (3.643)
Date FE	No	Yes	Yes
Site \times Time Trend	No	No	Yes
Observations	312	312	312

Notes: Heteroskedasticity-robust standard errors in parentheses. Observations from 9/1/2014 - 9/1/2016 at the weekly level. All columns include main effects for Ticketmaster and SeatGeek. Column 1 includes an indicator for after the experiment, *Post*.