Price Salience and Product Choice

Tom Blake, *eBay*
Sarah Moshary, *University of Pennsylvania*
Kane Sweeney, *Uber*
Steve Tadelis, *University of California, Berkeley*

April 11, 2017

PRELIMINARY AND INCOMPLETE

Abstract

We study the effect of price salience on product choice along two dimensions: whether a good is purchased and, conditional on purchase, the kind of good purchased. Consistent with our theoretical predictions, we find that making the full purchase price salient to consumers reduces both the quality and quantity of goods purchased. The effect of salience on quality accounts for at least one-third of the overall revenue decline.

1 Introduction

Textbook models of consumer choice assume that economic agents are fully aware of fees and taxes. Consumption decisions are therefore based on true final prices. As a consequence, any change in the final prices of goods, due to a change in base prices, fees, or tax rates, results in the same change in consumer choices. Recent influential work, however, offers evidence that challenges this assumption. Examples include Chetty et al. (2009), who document that tax salience profoundly affects consumers' decisions to purchase personal care goods in grocery stores; Finkelstein (2009), who finds that local governments exploit salience limitations to raise fees on toll roads; and Hossain and Morgan (2006), who find that eBay buyers respond more to list price than to shipping cost.[1] In this paper, we employ a large-scale field experiment to show that the effect of price salience on what consumers purchase can be just as important as the effect on whether they purchase.

Consider the example of a percentage fee levied on all goods. Price salience affects the consumer along two margins. First, increasing salience makes all goods appear more expensive, impacting the extensive margin and resulting in a higher probability that the consumer chooses not to buy any good. Second, because a percentage fee amounts to a larger fee level for more expensive goods, salience changes the perceived price-quality tradeoff for the goods in the consumer's choice set. As a result, higher salience impacts the intensive margin and encourages the consumer to substitute towards cheaper goods. The contribution of our paper is to offer a more complete analysis of the effect of price salience on consumers' choices by quantifying the importance of both margins.[2]

We begin our analysis with a simple model that illustrates the impact of price salience on consumption choices. The model demonstrates that if prices are made more salient—i.e., fees are listed upfront—then consumers are not only less likely to purchase any good, but conditional on purchasing, they

[1]See also Brynjolfsson and Smith (2001) and Sullivan (2017), who considers resort fees and hotel rates. In a related vein exploring choices that are more demanding on cognition, Allcott and Taubinsky (2015) find that consumers underestimate the cost savings from choosing energy efficient lightbulbs, and Abaluck and Gruber (2011) find that elders place more weight on medical plan premiums than on expected out-of-pocket costs.

[2]In 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.

purchase lower quality goods.

We take these predictions to data generated from a large-scale field experiment conducted by StubHub.com, the leading online secondary ticket marketplace. Until the experiment in 2015, the platform employed an Upfront Fee strategy, where the site showed consumers the final price from their very first viewing of ticket inventory. This final price included all ticket fees and taxes. The platform then experimented with a Back-end Fee strategy, where fees (such as shipping and handling) were shown only after consumers had selected a particular ticket and proceeded to the checkout page.

StubHub randomly selected 50% of users for the Upfront Fee experience (UF), while the other 50% had to make initial tickets selections based solely on the seller's listing price. This Back-end Fee (BF) group saw the fee-inclusive price only at the final stage of the checkout process. This experiment provides exogenous variation in fee salience in a setting where we can collect detailed data on consumer choice that includes the actual choice sets, signals of purchase intent (e.g., product selection and clicks towards checkout), and final purchase choices. These rich data allow us to infer the effect of salience on both the extensive and intensive margins of product choice. Our empirical results bear out the model's predictions: price obfuscation distorts both quality and quantity decisions. The intensive margin accounts for approximately one-third of the increase in revenue raised from fees.

Analysis of the Clickstreatm data suggests that Back-end Fees play on consumer misinformation. Upfront Fee users are more likely to exit before exploring any ticket, while Back-end Fee users differentially exit at checkout, when they first see the fee. Further, Back-end Fee users go back to examine other listings more often than their Upfront Fee counterparts. They are more likely to go back multiple times, which suggests Back-end Fees make price comparisons difficult. Back-end Fees affect experienced users, although on a smaller scale, which is consistent with consumers facing optimization costs, even when they anticipate a fee.

We also examine whether salience is more or less pronounced for big-ticket items. If consumers employ heuristics, they may not respond strongly to proportional fees for the sort of low-cost items studied in the existing literature (such as drug store beauty aids) but react more extremely to those same fees when levied on relatively costly products (such as $300 playoff tickets). Understanding when and where salience looms large is crucial to crafting both government tax policy and firm pricing strategies. Our data contravene this hypothesis: when hit ex post by an obfuscated fee, consumers are less likely to exit for higher priced tickets.

As a robustness test, we present evidence on price salience from an earlier experiment at StubHub performed in 2012, when the default user experience was BF. This second experiment randomized at the event-, rather than cookie-, level, and therefore suffers from a separate set of vulnerabilities. Reassuringly, the results are broadly consistent with our 2015 findings.

The next section presents a simple model of price salience and derives empirical implications. Section 3 discusses the experiment run at StubHub.com, as well as the data used in the analysis. Section 4 describes robustness checks on the randomization, while section 5 provides results. Section 6 contains evidence on mechanisms. Robustness checks based on the second experiment are presented in section 7. Section 8 concludes.

2 A Model of Consumer Choice with Limited Fee Salience

We consider consumers who decide whether and which products to buy under two regimes. Under the first regime, which we call Upfront Fees (UF), the final purchase price including all fees is shown to consumers when they browse through the set of available products. In the second, which we call Back-End Fees (BF), only list prices are shown to consumers when they browse through available products, and fees are revealed only after a particular product is selected for purchase.

First, we consider consumer i's choice when she observes fees upfront. She is presented with a set of available tickets J, where her utility v_{ij} from ticket $j \in J$ depends on its price p_j and quality δ_j in the following fashion:[3]

$$v_{ij} = \delta_j - \alpha_i p_j.$$

Consumer i's willingness to trade off quality for money is captured by α_i. For convenience, let $0 \in J$ denote the outside option, with $\delta_0 = p_0 = 0$. The figure below illustrate her optimization problem: the left panel displays the supply of tickets available to the consumer, while the center panel illustrates an indifference curve. The consumer then chooses the ticket from the supplied set on her highest indifference curve, indicated by the tangency on the right.

[3]Quality is a function of section, row and delivery option; e.g. instant download, FedEx, etc.

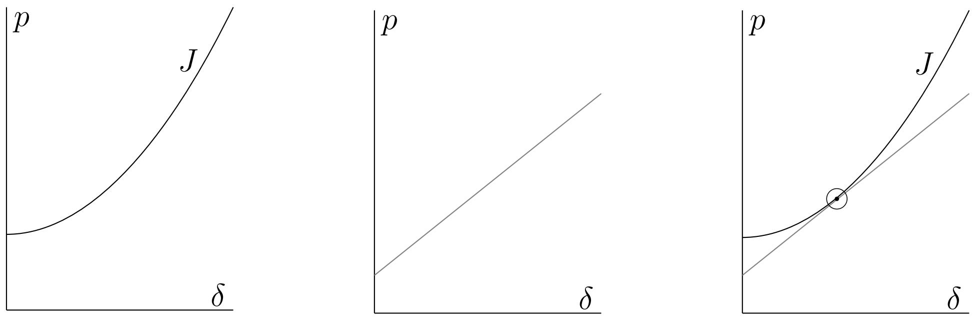

supply indifference choice

Following Finkelstein (2009), we model consumer optimization with Back-end Fees differently. Her choice now depends on the *perceived* price of ticket j, $\tilde{p}_j$, rather than its true final price (following the notation in Finkelstein 2009). The consumer then selects $j \in J$ to solve the following optimization problem:

$$\max_{j \in J} \tilde{v}_{ij} = \max_{j \in J} \delta_j - \alpha_i \tilde{p}_j$$

where the perceived price of not purchasing a ticket is also zero, $\tilde{p}_0 = p_0 = 0$. The established view on price salience implies that $\tilde{p}_j \leq p_j$. That is, when fees are obfuscated, prices appear lower to consumers.

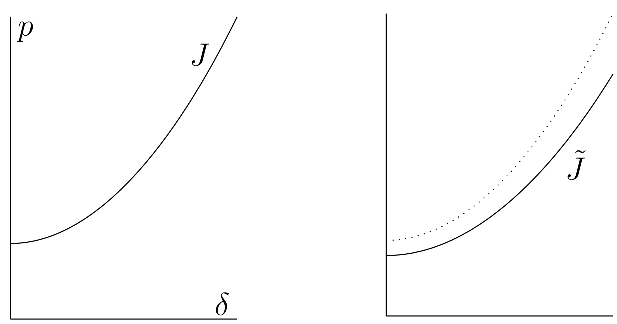

actual perceived

This simple model suggests that fee obfuscation has two effects on consumer behavior:

(1) Quantity Effect: *Under the Back-end Fee treatment, a consumer is more likely to purchase.*

This prediction is consistent with the existing literature: when fees are made salient, the likelihood of purchase declines. This model precludes at least two possibilities about the effects of salience: first, if consumers anticipate fees (or hold unbiased beliefs about fees) then our assumption that perceived prices are lower than actual prices might be faulty. 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.

On the other hand, if actual prices are higher than perceived prices and the difference is *increasing* in actual price, then this baseline 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 a proportional fee or tax.

(2) Quality Upgrade Effect: *If* $\tilde{p}_j - p_j \geq 0$ *and* $\tilde{p}_j - p_j$ *is increasing in* δ_j*, consumers buy higher quality tickets.*

Conditional on purchasing, consumers upgrade to higher quality tickets under Back-end Fees. They 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, Loney and Kroft 2009). Indeed, the log-log demand specification favored by earlier work leaves no scope for upgrades.

The Quality Upgrade Effect implies that consumer i buys a more expensive ticket. The increase in price, conditional on purchase, is illustrated in the following figure:

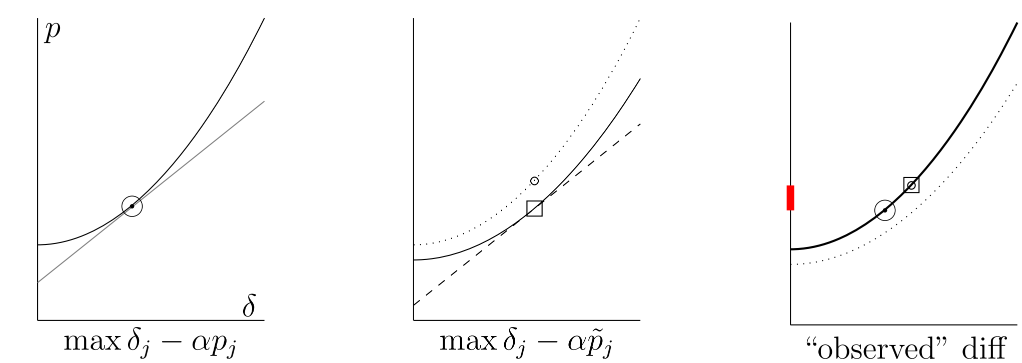

The Quality Upgrade Effect emphasizes how identification strategies must respect the impact of salience on quality choice. Consider the alcohol sales analysis of Chetty, Looney, and Kroft (2009). They compare an excise (lump sum) tax to a sales (percentage) tax. The excise tax should arguably not effect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

The next section elaborates our 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 examine an experiment in price salience on StubHub.com, a platform for secondary market ticket sales. Between 2013 and 2015, the platform showed all fees upfront. That meant the initial price a consumer saw when browsing ticket inventory was the final checkout price. In 2015, the firm ran an experiment during the final two weeks in August (August 19th - 31st). Treated consumers were shown ticket prices without the additional fees; these fees were only added at the checkout page, much like taxes added at the register of a grocery store.[4] We refer to this as Back-end Fees. The platform employs a non-linear fee structure: the buyer fee is 15% of the ticket price plus shipping and handling, if applicable. The platform also charges seller fees which peak at 15%.

The experimental condition was assigned at the cookie-level, which identifies a browser on a computer. 50% of 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 Upfront Fee 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, users in the test Back-end Fee group were shown only the base price when they perused available listings. Once a user in the Back-end Fee group selected a ticket, they were taken to a ticket details page,

[4]Other ticket platforms, including Ticketmaster, employ a similar pricing scheme, where fees are only included at the final stage of the transaction.

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 proceed to checkout or abandon the purchase.

We exploit randomization of the BF experience to estimate the quantity effect described in section 2: the quantity effect is the difference in the likelihood of purchase between UF and BF users. Let Q_i be the likelihood consumer i purchases and T_i be an indicator for Back-end Fees, then

$$\text{Quantity Effect} = E[Q_i|T_i = 1] - E[Q_i|T_i = 0]. \tag{1}$$

We implement (1) using an OLS framework. 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:

$$Q_i = \alpha + \beta \cdot T_i + \epsilon_i \tag{2}$$

The coefficient β represents the difference in the levels of Q_i for Back-end compared to the Upfront Fee users. To protect business-sensitive information, however, we report estimates of $\frac{\beta}{\alpha}$, which is the percent change for Back-end Fee users.

Unfortunately, it is difficult to test whether the Quality Upgrade Effect is positive through random assignment of the Back-end Fee experience. Let P_i be the purchase price of the ticket that consumer i selects, conditional on i making a purchase. Let Q_{i0} be an indicator for whether consumer i purchases a ticket when he observes fees upfront ($T_i = 0$) and Q_{i1} when he observes them at the back-end ($T_i = 1$). Using the potential outcomes notation, the quality upgrade formulation is:

$$\text{Quality Upgrade Effect} = E[P_i|Q_{i0} = 1, T_i = 1] - E[P_i|Q_{i0} = 1, T_i = 0]. \tag{3}$$

The econometrician cannot observe the first term, which is the average purchase price among Back-end Fee users who *would have* bought under Upfront Fees. Instead, we observe the change in the average price, conditional on purchasing.

$$\begin{aligned} \Delta\ \text{Purchase Price} = E[P_i|Q_{i,1} = 1, T_i = 1] - E[P_i|Q_{i,0} = 1, T_i = 0] \\ \geq \text{Quality Upgrade Effect} \end{aligned} \tag{4}$$

This change in average purchase price constitutes a lower bound for the quality upgrade effect under a monotonicity assumption ($Pr\{Q_{i,1} = 1|Q_{i,0} = 1\} = 1$). It combines two separate effects: first, the quality upgrade effect, which encourages consumers to purchase more expensive tickets than they would otherwise and second, a change in the marginal consumer, as more consumers purchase tickets. The latter depresses average ticket prices in the Back-end Fee group if marginal consumers buy the cheapest tickets. Indeed, marginal consumers ought to be more price sensitive than infra-marginal consumers, and therefore likely to buy inexpensive tickets. The change in average purchase price (4) is therefore a lower bound for the quality upgrade effect. We estimate (4) using (2) with price on the left-hand side.

We use conditional probability to derive an upper bound for the quality upgrade effect. Essentially, the bound sets the average price for marginal consumers at zero (that is, users who buy under BF but abstain under UF get tickets for free under the BF treatment). We relegate the details to the appendix.

$$\begin{aligned}\text{Quality Upgrade Effect} \leq \frac{Pr\{Q_{i,1} = 1\}}{Pr\{Q_{i,0} = 1\}} \\ \cdot\Big(\Delta \text{ Purchase Price } + E[P_i|Q_{i,0} = 1, T_i = 0] \cdot (1 - Pr\{Q_{i,0} = 1\})\Big)\end{aligned} \tag{5}$$

We note that the change in average purchase price is inherently interesting in this setting, as it constitutes a change in platform revenue. We decompose expected revenue using conditional probability as:

$$E[R_i] = E[P_i|Q_i = 1] \cdot Pr\{Q_i = 1\} = E[P_i|Q_i = 1] \cdot E[Q_i].$$

So that the change in revenue from treatment is

$$\Delta E[R_i] = \underbrace{\Delta E[P_i|Q_i = 1]}_{\Delta \text{ Purchase Price}} \cdot E[Q_i] + \underbrace{\Delta E[Q_i]}_{\substack{\text{Quantity} \\ \text{Effect}}} \cdot E[P_i|Q_i = 1].$$

4 Randomization Check

In total, the experiment included several million visitors who frequented the site over ten days. To check randomization, we test whether we can reject a 50% treatment assignment probability. Results are presented in table 1.

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, and second, users on a particular browser-operating system combination were also skewed to the treatment group. Once we eliminate these two groups, we can no longer reject a 50% assignment at the 1% level.[5] We therefore exclude these users in our main analysis. Although the probability of treatment remains slightly higher than 50%, the difference is economically insignificant.

Table 1: Treatment Assignment

Sample	% in Niether Back-end nor Upfront Fees	% 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

As a robustness check on randomization, we test whether UF and BF users share similar observable characteristics. Unfortunately, since treatment was assigned before users are required to log-in, the set of observables is fairly limited. As an example, even if a user has visited the site before, we do not know their purchase history if they do not log into the site during the experiment and they have also cleared their cookies. However, we are able to measure site visits since the last cookie-reset, which we use to measure experience. We employ this as a left-hand side variable in specification (1). Table 2 reports the regression results, which show 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 be mac users. We devote the remainder of the paper to a detailed analysis of the experiment.

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

Table 2: Covariate Balance

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

5 Effect of Salience on Revenue

Theory indicates that obfuscation should encourage marginal consumers to switch from buying nothing to buying something, and also encourage consumers to switch from purchasing lower to higher quality tickets. Table 3 shows the composite effect on revenue of the price salience treatment. Consumers identified with cookies in the Back-end Fee group, where fees are shrouded, 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).

Unfortunately, quantifying salience is difficult so it is hard to compare this effect to the Chetty, Looney and Kroft (2009) benchmark. (While the change in user experience in the StubHub experiment is similar in spirit to their experiment of adding fees to supermarket shelf prices, it is hard to measure 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.[6] Our findings, detailed below, suggest that upgrades induce a larger salience effect in this setting.

5.1 Quantity Effect

We first examine the quantity impact, defined in (1). The third row of table 3 shows that price obfuscation increased the transaction rate over the full course of the experiment. We find that consumers in the Back-end Fee group are 13% more likely to purchase a ticket during a visit. Fees average roughly 15% of ticket prices, suggesting a salience elasticity of 0.87, which is similar

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

Table 3: 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)	_

to the elasticity of 1.1 found in Chetty, Loony & Kroft. The 10-day elasticity is larger than the session elasticity, suggesting that the long-run effects of salience may be even larger.

Our simple model of consumer choice abstracts from the consumer's decision of how many seats to buy, conditional on purchase. The model maps to a world where consumers need a fixed number of seats and do not consider purchasing any more or fewer, unless they choose the outside option. In reality, of course, consumers might decide to 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.com. 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 fewer tickets.

5.2 Quality Upgrade Effect

The second column of Table 3 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: Back-end Fee users spend 5.42% more than their Upfront Fee counterparts. From the platform's perspective, this implies that the effect of salience on their bottom line is substantially larger (on the order of 30%) than suggested in the earlier literature, which did not consider product choice. Note that the number of seats declines slightly, so that the change in the average purchase price per seat is even greater (5.73%).

Figure 1: Likelihood of Purchase by Row for BF versus UF users

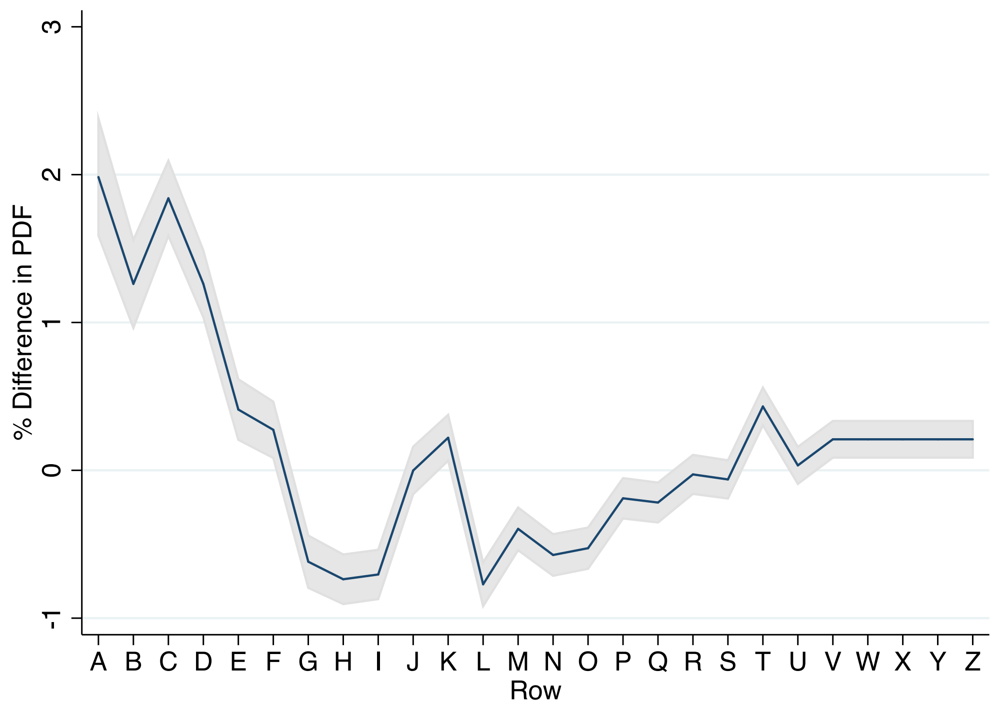

We interpret this as evidence of an upgrade effect, where shrouding fees leads consumers to buy higher quality tickets than they would otherwise purchase. Our model of consumer choice indicates the change in the average purchase price constitutes a lower bound for the upgrade effect – and while smaller than the quantity effect, even the lower bound is economically meaningful. Our calculation of the upper bound (6) is 16.52%, suggesting that the Quality Upgrade Effect may hit just as hard as the Quantity Effect.

While quality is difficult to measure directly, we also examine whether Back-end Fee users buy seats closer to the stage. Rows are often labeled using letters, where letters earlier in the alphabet correspond to a better view. Of course, seat, row, and section numbering schemes vary substantially across venues, so this is only a proxy for quality. Conditional on purchasing a ticket, we separately calculate the probability that a BF and UF user purchases a seat in each row. Figure (1) graphs the relatively probability (the ratio of the probability mass functions), along with 95% confidence intervals (which are calculated point-wise). Back-end Fee users are more relatively more likely to purchase seats in row A, and the likelihood declines for rows later in the alphabet. These trends provide further evidence of the quality upgrade effect.

Figure 2: Transaction Per Ticket Price Distribution

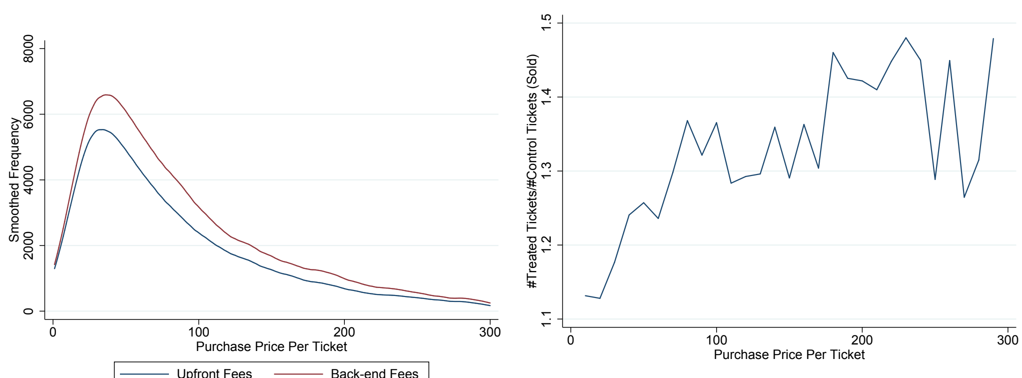

(a) Kernel Smoothed Histogram

(b) Ratio of Back-end to Upfront Fee Purchase Counts

6 Mechanisms

6.1 Heterogeneity across Price Points

We first consider the heterogeneity of response by price range. If consumers employ heuristics, then we might expect the effect of price salience to attenuate at higher price points. That is, consumers should pay closer attention to fees when the fee level is higher.

Panel (a) in Figure 2a shows the overlapping Back-end and Upfront Fee distributions (kernel smoothed) of ticket prices for all purchased tickets with prices less than \$300 during the experiment. The response is surprisingly consistent across all price points. But it is not completely uniform, as is clear in panel (b) of Figure 2b.

Table 4: Relative Purchase Probability by Ticket Price

Ticket Price	Relative Purchase Probability by Back-end versus Upfront Fee Users
< \$20	1.02
\$20 - \$100	1.13
\$100 - \$200	1.23
\$200+	1.27

Table 4 shows that the difference in purchase probability between treatment and control is barely noticeable for tickets priced below \$20. However, the difference becomes significantly larger for tickets between \$20 - \$100, and even more pronounced for tickets above \$100. These results are hard to reconcile with the rational-inattention story.

Table 5: Purchase Funnel Behavior by Fee Salience

	Percentage Click Through from Prior Page			Average Ticket Price		
	BF	UF	% Difference	BF	UF	% Difference
Event Page	_	_	_	$1.00	$0.84	18.73%
Ticket Details	63.56%	54.82%	15.94%	$0.86	$0.78	10.16%
	(0.01)	(0.01)	(0.03)			
Review & Submit	_	_	_	$0.56	$0.52	7.44%
Purchase	38.08%	54.23%	-29.78%	$0.42	$0.39	6.57%
	(0.01)	(0.02)	(0.03)			

6.2 Misinformation

We leverage StubHub.com'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 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, this theory has an implication about where (in the purchase funnel) Back-end and Upfront Fee users should differentially exit.

To buy a ticket, a user goes down the StubHub 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;[7] (3) the consumer reaches the checkout page where a final purchase decision is made; (4) the purchase confirmation page completes the process.[8] BF users are shown lower prices than their control 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 smaller than the latter, then Back-end Fees increase quantity sold.

The right panel of Table 5 shows the absolute and relative rate of UF and BF user arrivals at each step in this process. Consistent with misinformation, Back-end Fee users are 16% more likely to select tickets (transition from stage 1 to 2) than users who see fees upfront. 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 large for BF users: 30%. Interestingly, the difference in drop off rates between stages 2 and 3 is also statistically significant (although smaller in magnitude than the others). We hypothesize that this difference results from selection; the users who differentially attrit due to price salience are more price sensitive.

The left panel of Table 5 examines the quality response at each step in the purchase funnel for a subset of events. The average price of tickets under

[7] the log-in page (optional; bypassed if the consumer is already logged into his account)

[8] Of course, many searches are non-linear Blake et al. (2016), where consumers examine multiple event pages. Back-end Fee users might even return to the search stage (1) once they see the additional fees leveed at stage (4).

consideration declines at each step in the funnel. Upfront Fee experience prices are lower than their Back-end counterparts, but the difference generally narrows as users move closer to purchase. At the point where fees are revealed, the gap is 7% compared to an initial difference of 19%. BF users, who see no fees, are more likely to contemplate expensive tickets. When fees are revealed, more of the (surprised) BF users exit than the UF users who see no change in their expected outlay.

One important question, from both the firm and policy 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 contrasts to a model where consumers anticipate a fee, but do not know the exact level. 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 3) separately by 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). Examining responses across experience groups, however, hints at how learning might work in this setting.

To measure experience, we calculate the number of visits each cookie has made to StubHub.com prior to the experiment. A 2006 ComScore study finds that 31% of users clear their cookies within 30 days, so we interpret this as a short-term measure of experience.[9] Unfortunately, we cannot exploit information about logged-in users (such as number of past transactions), since 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, since it is a prerequisite to purchase. Importantly, however, our measure captures the most recent interactions with StubHub, which are likely to be the most relevant for a users' knowledge of the site.

We hypothesize that frequent StubHub.com users ought to be aware of fees and therefore less sensitive to price 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 6 shows that the treatment effect is smaller for cookies with at least 10 site visits: the revenue effect is 15% compared to

[9]https://www.comscore.com/Insights/Blog/When-the-Cookie-Crumbles

Table 6: Salience by User Experience

	% Difference		
	New User	Low Experience	High Experience
User 10-day Revenue	21.52%	21.80%	15.09%
	(0.02)	(0.02)	(0.04)
Propensity to Purchase at Least Once	15.33%	13.68%	10.19%
	(0.007)	(0.01)	(0.02)
# Transactions within 10 Days	14.33%	13.53%	8.81%
	(0.01)	(0.01)	(0.03)
% Sample	67%	27%	6%

21%. These results suggest that salience effects ought to be most important in markets where consumers purchase infrequently (for example, real estate or automobile markets). However, price obfuscation still generates substantial revenue, which indicates only limited consumer learning – even experienced users seem to be caught off-guard by fees.

We also examine user churn. If obfuscation preys on misinformation, then marginal BF consumers should be unlikely to return to the site. These are users who would not purchase if shown fees upfront, so that after they see fees for the first time, marginal consumers users should eschew StubHub. Unfortunately, we cannot identify marginal consumers. Nor can we 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 window. As shown in table 3, BF users are 3.5% less likely to return to the site within 12 months. However, this difference potentially confounds multiple treatments; Back-end Fee users may learn about the platform fees when they make a purchase, but they may also learn about StubHub's reliability, speed, quality, etc. We next compare the likelihood of return for consumers who were logged in to the StubHub website *before* the experiment. We can track these users future purchases, 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).

6.3 Search Frictions

Obfuscating fees may lead consumers – particularly first-time visitors – to spend more money on the platform than they would otherwise. When fees are revealed, BF consumers are already at check-out with their tickets. They must choose to go back to the event page if they wish to re-optimize and purchase other seats. Figure 3 shows the average number of tickets viewed for BF and UF users. BF cookies are 56% more likely to view multiple ticket listing compared to their UF counterparts. Table 7 shows that BF users view cheaper tickets upon their return (6 percentage points lower). In contrast, UF users (who are less likely to return) go back for relatively more expensive tickets.

Figure 3: Number of Listings Viewed by Fee Salience

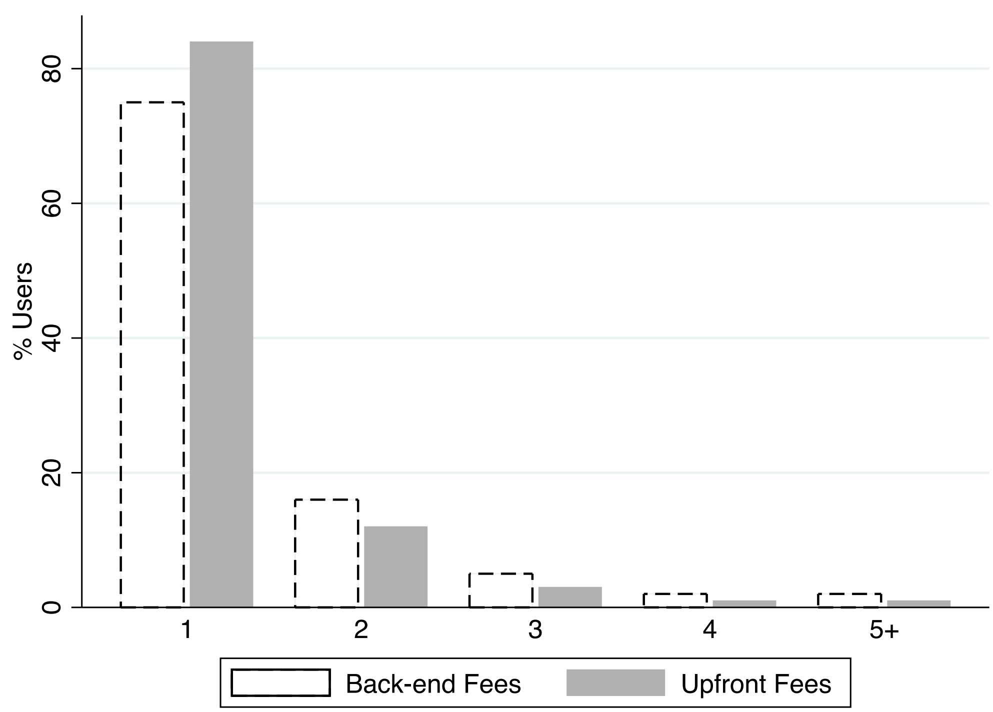

Pearsons Chi-square of 6700 rejects hypothesis that distribution over rows is same in test and control. (p-value of 0.000)

Back-end Fee users are twice as likely to view 3 or more listings than their Upfront Fee 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 already for their initial selection, but they must calculate the StubHub fee for each new ticket they consider. If this calculation cost is high, consumers might choose to go down the funnel multiple times, so that StubHub reveals the final price, rather than compute the fees themselves. Obfuscation as a search friction is consistent 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 consistent with Ellison and Fisher Ellison (2009), who find that firms endogenously create such frictions to soften price competition.

Table 7: 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)	(0.5)	(0.6)

6.4 Endowment Effect

Finally, we investigate whether Back-end Fees create attachment bias. In a seminal paper, Khaneman and Tversky (1979) suggest that consumers value objects differently when they feel ownership over the good. In this case, consumers who don't anticipate fees may put tickets in their "cart" and be loathe to part with these tickets later, even when fees are revealed. By hiding fees, the platform changes the consumer's utility function at the purchase juncture.

The endowment effect logic works as follows: BF users don't anticipate fees. Once they see fees, some decide to go back. Who decides to return given that search is costly? The folks who are most price sensitive. These consumers should choose relatively cheaper seats compared to the average UF cookie. But the seats they do choose are actually more expensive than the average AIP price. Table 7 shows that returning BF users interact with 2% more expensive tickets than UF users. This suggests that the BF folks who go back put a higher premium on quality, consistent with an endowment effect.

6.5 Competition with Other Platforms

We have shown that Upfront Fees reduce the number of users who buy tickets on the site by 14%. These marginal consumers might exit the market altogether or they may purchase tickets through a rival platform. While fee salience remains essential to the platform's bottom line in either case, understanding where marginal consumers go has important welfare implications. As an example, if all sellers multi-home, then consumers might buy the same exact tickets on Ticketmaster that they would have under Back-end Fees at StubHub. Obfuscation would then have only limited efficiency consequences (through the product selection margin for users who remain on StubHub). On the other hand, if consumers who leave StubHub under Upfront Fees exit the market, then the change in consumer surplus could be much larger.

To investigate the effect of StubHub's switch to Back-end Fees, we employ data from GoogleTrends on queries for its main competitors: Ticketmaster and SeatGeek. Both sites act as a secondary market for tickets, with Ticketmaster serving as the primary market for certain sporting and music events. Google provides data on weekly query volume for these sites, but normalizes the data separately for each platform (by dividing by the site's peak over 2012-2017). Figure 4 shows the evolution of queries for 2015 and 2016. A Chow test indicates a break in August 2015 for Ticketmaster and StubHub, but not for SeatGeek. The lift in StubHub queries suggests that consumers deterred by Upfront Fees at StubHub do not turn to alternative sites.

6.6 Seller Responses

In this section, we provide evidence on seller responses to fee salience. Our main analyses consider impacts on consumer behavior, but 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.

As a first step, we consider whether sellers alter their pricing in response to fee obfuscation. Sellers choose a base price for each listing, on which StubHub imposes a percentage fee. Under Back-end Fees, StubHub displays

Figure 4: Google Queries for Competing Ticket Resale Platforms

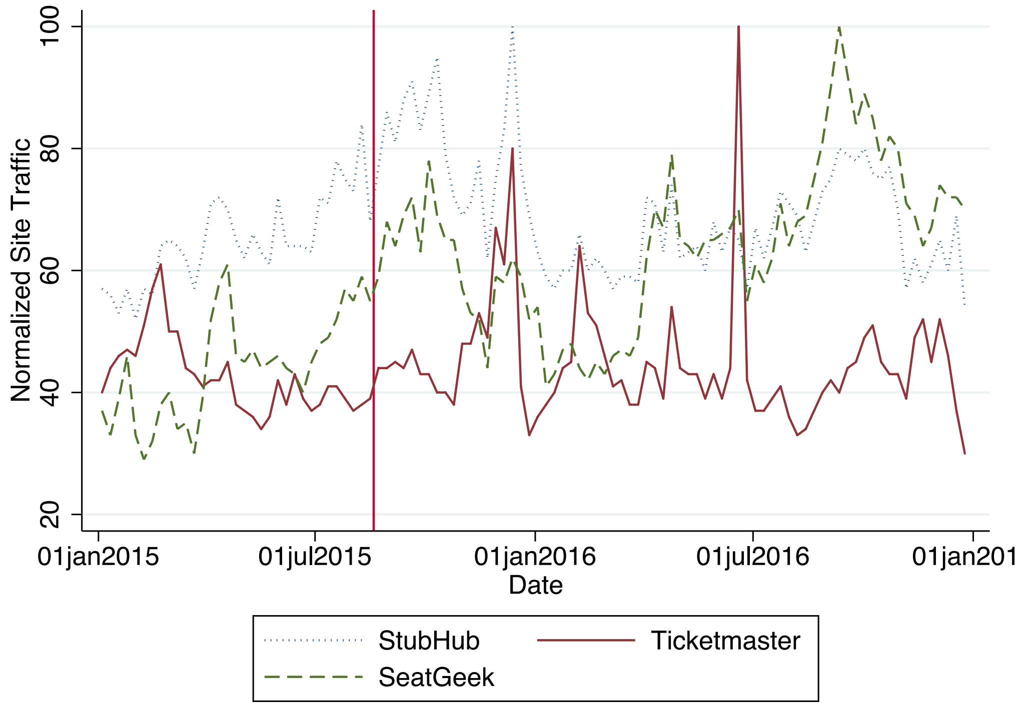

the seller base price to potential buyers until they reach the Review & Submit page, when the user is shown total price (seller price + shipping + buyer fees). Under Upfront Fees, StubHub displays the total price from the get-go. Sellers may alter their prices depending on the buyer experience to exploit the limited fee salience. As an example, an extensive literature in marketing documents the appeal of round number pricing (amounts that end in zeros or nines).[10] We examine whether sellers are more likely to set base prices at round numbers after the switch to Back-end Fees.

As shown in figure 5, the share of listings that are round increases by approximately 5 percentage points following the switch to Back-end Fees. Sellers seem to respond to the change in the buyer's experience, which is consistent with Ellison and Fisher Ellison (2009).

[10]For example Monroe (1973) or more recently Backus et al. (2015).

Figure 5: Percent of Listings that use Round Numbers

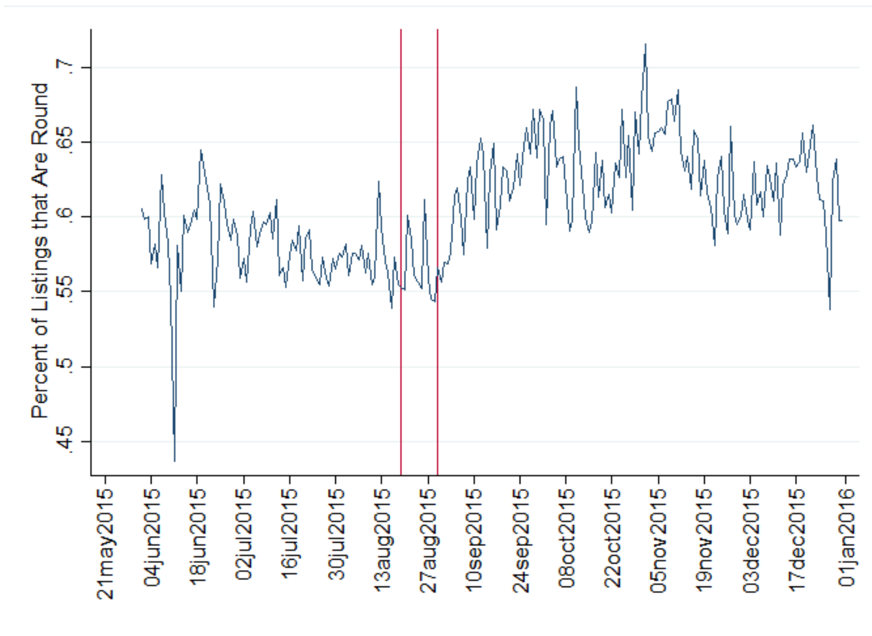

7 A second experiment: randomization at the event-level

In the 2015 experiment, feel salience was randomized across cookies. Back-end Fee and Upfront Fee users had the same StubHub experience, but for the addition of fees in the latter's search results. In a second experiment at StubHub, fee salience was randomized at the event level. Randomization at the event level presents distinct challenges, but offers a nice robustness check for the 2015 experiment.

The uniqueness of StubHub 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 example, Back-end Fee users will tend to buy early in the 2015 experiment. If BF users buy up all the inventory, then UF users can no longer purchase a ticket. 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. Blake and Coey (2014) discuss this challenge on eBay.com. Fortunately, the 2012 experiment does not suffer from the same contamination concern because all tickets for a particular event share the same treatment status. In the example above, there would be no difference in sales across Back-end and Upfront Fee games, only a difference in the timing of purchases.

A second challenge that the 2012 experiment addresses is multi-homing. In the 2015 experiment, we sort users into the Back-end or Upfront Fee group the first time they touch an event page on StubHub.com during the experiment period. StubHub.com employs cookies to track users, so that the user remains in the appropriate group throughout the trial. However, the cookie does not follow the user if he were to visit StubHub.com on a second computer or on a mobile device. Instead, the user would be re-randomized into the BF or UF group. Multi-homing is particularly problematic if its incidence depends on initial treatment assignment. As an example, if Upfront Fee users – upon seeing higher initial prices – delay their purchases and revisit StubHub.com on a second device, then treatment would be artificially correlated with purchasing. In the 2012 experiment, tickets to each event retain their treatment status regardless of the device that consumers employ. If a Red Bulls vs Revolution match shows the fee-inclusive price on their personal laptop, it also shows the fee-inclusive price on their work Desktop.

Finally, randomization at the event level provides insight into general equilibrium effects. When StubHub.com alters the consumer's experience, it potentially alters sellers' incentives. As an 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 on StubHub.com when price is obfuscated for all users? Instead, the econometrician only observes what happens on StubHub.com when price is shrouded for 50% of users. The 2012 experiment provides insight into the importance of these GE effects, since a ticket-seller for a particular match faces an entirely Back-end or Upfront Fee audience, but not a mix of both.

In the 2012 experiment, 33 out of 99 Major League Soccer Games were randomly selected for Upfront Fee. Prices for tickets to these games included fees, even from the initial event page. For the remaining 66 matches, fees were added only in the final checkout stage of the purchase funnel.

Table 8: 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: Standard errors clustered at the event level.

The results from the 2012 MLS experiment, displayed in table 8 , confirm our 2015 findings: fee salience reduces revenue substantially. Consumers are 13% less likely to buy tickets to an Upfront Fee match (note that fees were approximately 10% in 2012). 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 Back-end Fee games. Unfortunately, tests based on purchase prices are underpowered because of the high sampling variance across matches. To control for the unobserved popularity of each match, 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 user's choice set (StubHub's entire inventory at the time of purchase for the match in question). On average, consumers buy from a 12% lower quantiles for UF compared to BF games. Figure 6 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. Experimentation at the event-level suffers from a different kind of contamination bias. The chief concern is that consumers may substitute away from Upfront Fee matches (which appear more expensive) to Back-end Fee matches. The 2015 experimental design is not vulnerable to

Figure 6: Percentile of Choice Set Purchased in the 2012 Experiment

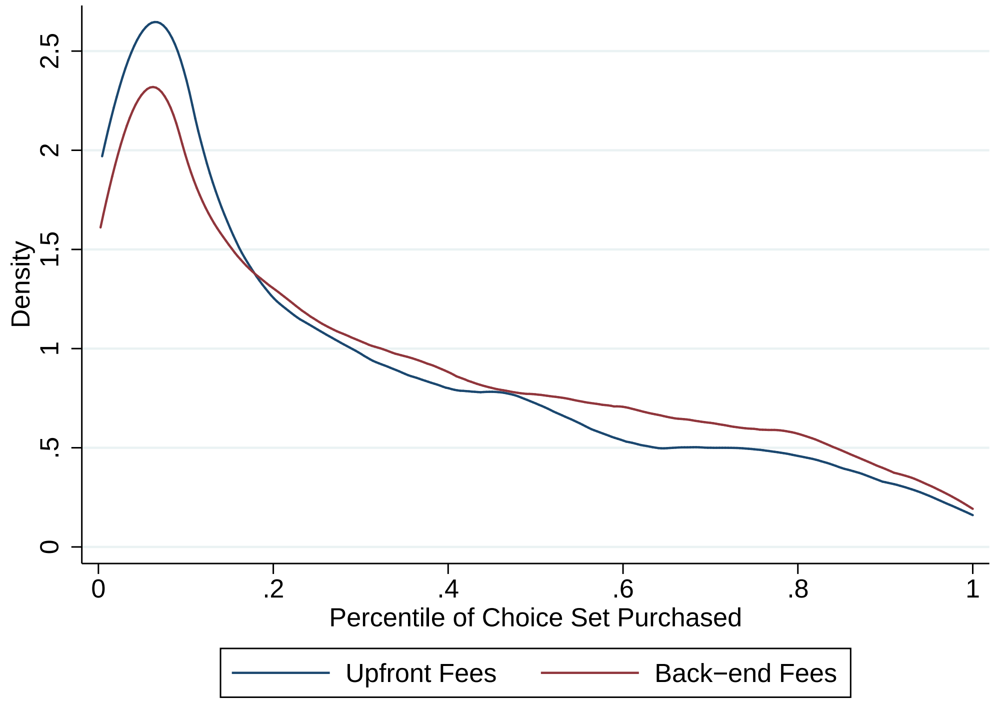

this type of contamination. The ability to execute two experimental designs is a nice advantage of the StubHub.com setting.

8 Conclusion

In a randomized control trial on StubHub.com, we find that shrouding buyer fees substantially increases total revenue. 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, which accounts for at least 30% of the revenue bump. The latter suggests consumers upgrade to higher quality products when they observe lower prices. We find consumers who are shown fees upfront drop-off early in the purchase funnel, while those shown fees later exit later (only once the site displays total prices). The hazard rates are consistent with consumer misinformation.

The effect of salience abates only slightly in a comparison of experienced users. Even users who choose to conduct a second search (after observing the total price for their initial selection) select more expensive goods when fees are less salient. This evidence suggests that obfuscation is not a one-off phenomenon, which becomes irrelevant as consumers learn about the sales environment. To the contrary, it indicates that site design can have a profound impact on consumer behavior.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *The American Economic Review*, June 2011, *101* (4), 1180–1210.

Allcott, Hunt and Dmitry Taubinsky, "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *The American Economic Review*, 2015, *105* (8), 2510–2538.

Backus, Matthew, Thomas Blake, and Steven Tadelis, "Cheap Talk, Round Numbers, and the Economics of Negotiation," *NBER Working Paper*, 2015.

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.

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.

Ellison, Glenn and Sarah Fisher Ellison, "Search, Obfuscation, and Price Elasticities On The Internet," *Econometrica, Vol. (, 2009),*, March 2009, *77* (2), 427–452.

Finkelstein, Amy, "E-Z Tax: Tax Salience and Tax Rates," *The Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.

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.

Monroe, Kent B., "Buyers' Subjective Perceptions of Price," *Journal of Marketing Research*, 1973, *10* (1), 70–80.

Sullivan, Mary W., "Economic Analysis of Hotel Resort Fees," *Economic Issues, Bureau of Economics Federal Trade Commission*, 2017.

An Upper Bound for Quality Upgrade Effect

$$\begin{aligned}
\Delta \text{ Purchase Price } &= E[P_i|Q_{i,1}=1,T_i=1] - E[P_i|Q_{i,0}=1,T_i=0] \\
\\
&= E[P_i|Q_{i,1}=1,T_i=1,Q_{i,0}=1]\cdot Pr\{Q_{i,0}=1|Q_{i,1}=1\} \\
&+ E[P_i|Q_{i,1}=1,T_i=1,Q_{i,0}=0]\cdot Pr\{Q_{i,0}=0|Q_{i,1}=1\} \\
&- E[P_i|Q_{i,0}=1,T_i=0] \\
\\
&= \text{Quality Upgrade Effect}\cdot\frac{Pr\{Q_{i,0}=1\}}{Pr\{Q_{i,1}=1\}} \\
&+ E[P_i|Q_{i,1}=1,T_i=1,Q_{i,0}=0]\cdot Pr\{Q_{i,0}=0|Q_{i,1}=1\} \\
&- E[P_i|Q_{i,0}=1,T_i=0]\cdot\left(1-\frac{Pr\{Q_{i,0}=1\}}{Pr\{Q_{i,1}=1\}}\right)
\end{aligned}$$

Rearranging, we have

$$\begin{aligned}
&\text{Quality Upgrade Effect} = \\
&\qquad\frac{Pr\{Q_{i,1}=1\}}{Pr\{Q_{i,0}=1\}}\cdot\Bigg(\Delta \text{ Purchase Price } + E[P_i|Q_{i,0}=1,T_i=0]\cdot Pr\{Q_{i,0}=0\} \\
&\qquad\qquad - E[P_i|Q_{i,1}=1,T_i=1,Q_{i,0}=0]\cdot(1-\frac{Pr\{Q_{i,0}=1\}}{Pr\{Q_{i,1}=1\}})\Bigg)
\end{aligned}$$

We construct an upper bound by dropping the last term, which is the expected purchase price for marginal consumers multiplied by a probability, and therefore weakly greater than zero:

$$\text{Quality Upgrade Effect} \leq \frac{Pr\{Q_{i,1}=1\}}{Pr\{Q_{i,0}=1\}}\cdot\Big(\Delta PP + E[P_i|Q_{i,0}=1,T_i=0]\cdot(1-Pr\{Q_{i,0}=1\})\Big) \tag{6}$$