Price Salience and Product Choice

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

February 27, 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 is approximately half the size of the effect on quantity.

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 purchase lower quality goods.

[1] 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.

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 provides results. Section 5 contains evidence on mechanisms. Robustness checks based on the second experiment are presented in section 6. Section 7 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 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.

A profit-maximizing firm cares ultimately about the effect of fee presentation on average revenue per consumer. Anticipating our experiment design, we estimate the revenue effect of fee revelation simply as the difference in average revenue across the BF and UF groups of consumers,

$$\Delta E[R_i] = E[R_i|T_i = 1] - E[R_i|T_i = 0]$$

where T_i is an indicator for whether a consumer i is treated (BF) and R_i is revenue from consumer i. However, this net change in revenue subsumes two separate effects: one on price and one on quantity. Let $Q_i \in \{0, 1\}$ be an indicator for whether consumer i purchases any product and P_i be the price of the product that i purchases. We decompose revenue using conditional probability as:

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

So that the change in revenue from treatment is

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

We employ the standard discrete choice model from Industrial Organization to provide predictions regarding the signs of these two effects. Let the utility of consumer i from purchasing product j be:

$$u_{ij} = v_j - \alpha \cdot p_j \cdot \theta(T_i) + \epsilon_{ij}$$

where v_j is the mean utility of product $j \in J$, p_j is product j's final price including fees, α is the shadow value of money, and $\theta(T_i) \leq 1$ captures the consumer's price salience, which depends on the treatment. In particular, the established view on price salience implies that $0 < \theta(T_i = 1) < \theta(T_i = 0) = 1$. That is, when fees appear in the Back End, they are less salient to consumers.

2.1 Quantity Effect

Under the standard logit assumptions on ϵ_{ij}, the probability consumer i purchases good j is:

$$q_{ij} = \frac{\exp\{v_j - \alpha p_j \cdot \theta(T_i)\}}{1 + \sum_{k \in J} \exp\{v_k - \alpha p_k \cdot \theta(T_i)\}}$$

The effect of the BF treatment on the expected likelihood of purchase, which we refer to as the *quantity effect*, is therefore,

$$\begin{aligned} \Delta E[Q_i] &= \Delta \left[\frac{\sum_{k \in J} \exp\{v_k - \alpha p_k \cdot \theta(T_i)\}}{1 + \sum_{k \in J} \exp\{v_k - \alpha p_k \cdot \theta(T_i)\}} \right] \\ &= -\alpha \cdot \Delta\theta \cdot (1 - E[Q_i]) \cdot E[R_i] > 0, \end{aligned}$$

where $\Delta\theta = \theta(T_i = 1) - \theta(T_i = 0) < 0$. Not surprisingly, the parametric logit model predicts that the BF treatment increase the likelihood of purchases. The magnitude of the quantity effect is increasing in price sensitivity α.

Rather than separately estimate the components of the quantity effect defined above, we estimate the quantity effect directly. We use the sample moments from the experiment that correspond to the treatment and control population counterparts. Because suppliers who sell tickets on Stubhub cannot price discriminate between BF and UF users, we need not worry that the two groups face different prices. Hence,

$$\Delta E[Q_i] = E[Q_i|T_i = 1] - E[Q_i|T_i = 0].$$

Importantly, we rely only on the assumption that $\Delta\theta < 0$ and need not lean on the extreme value assumption to estimate the quantity effect. Randomization permits us to test the model's sign prediction. As an example, if price obfuscation generates a 'disgust' factor, wherein adding fees at the end upsets consumers, then our estimated quantity effect would be negative, leading us to reject the standard price salience model.

2.2 Price Effect

The model also predicts that Back-end Fees increase the expected purchase price (conditional on purchase). We refer to this as the *price effect*, which is given by,

$$\begin{aligned}\Delta E[P_i|Q_i = 1] &= -\alpha \cdot \Delta\theta\Big(E[P_i^2|Q_i = 1] - E[P_i|Q_i = 1]^2\Big)\\ &= -\alpha \cdot \Delta\theta \cdot Var(P_i|Q_i = 1) > 0.\end{aligned}$$

The intuition underlying this result is simple: in a logit model, the conditional variance of price is high when there are goods of similar total utility at different price points. When a proportional fee is hidden, the utilities of expensive goods increases disproportionately, leading to substitution.[3] Again, we can test this implication directly by comparing the average purchase price in the BF and UF groups:

$$\Delta E[P_i|Q_i = 1] = E[P_i|T_i = 1] - E[P_i|T_i = 0]$$

Ex ante, it is unclear whether average prices should rise among Back-end compared to Upfront Fee purchasers. So long as perceived prices are weakly less than final prices, then there are two components to the price effect: (1)

[3]There is no such price effect, for example in a log-log demand specification.

more consumers will purchase at least one ticket and (2) consumers will purchase more expensive tickets than they would otherwise. (1) may actually depress average ticket prices in the Back-end Fee group if marginal consumers buy the cheapest tickets. Our simple discrete choice model ignores consumer heterogeneity, so there is no change in the identity of the marginal consumer.

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, so that the initial price that 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 treatment 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 where the control group were shown the total price (ticket cost plus fees and shipping charges). Users could thenproceed to checkout or abandon the purchase.

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

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

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

We estimate the effect of Back-end Fees on user's purchase decisions using a standard OLS regression model:

$$y_i = \alpha + \beta \cdot T_i + \epsilon_i \tag{1}$$

where y_i is the outcome for cookie i, such as a purchase dummy or total spent on the site, and T_i is a treatment indicator that takes a value of one if cookie i sees fees only at the end of the purchase funnel (BF). Since assignment to Back-end Fees is random, we do not include additional control variables. The coefficient β represents the difference in the levels of y_i for Back-end Fee compared to the Upfront Fee users. To protect business-sensitive information, we report estimates of $\frac{\beta}{\alpha}$, which is the % change for Back-end Fee users.

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

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

limited. As an example, even if a user has visited the site before, we do not know their purchase history if they never 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. Treatment and control 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.

Table 2: Covariate Balance

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

4 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. In our setting, obfuscating a 17% fee leads to a 21% revenue boost.

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 (1.03)
Propensity to Purchase at Least Once	14.10 (0.09)	_
# Transactions within 10 Days	13.24 (0.88)	-0.90 (0.37)
Cookie Session Revenue	18.96 (1.27)	5.61 (0.96)
Cookie Session Propensity to Purchase	12.43 (0.6)	_
Cookie-Session # Transactions	11.76 (0.54)	-0.66 (0.38)

4.1 Quantity Effect of Salience

We next examine the quantity impact separately. 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 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. In Section 5.1, we explore consumer behavior in more detail by looking at the rate that users engage with each phase of the purchase funnel.

4.2 Price Effect of Salience

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 the price effect: Back-end Fee users spend 5.42% more than their Upfront Fee counterparts. We interpret this as an upgrade effect, where shrouding fees leads consumers to buy higher quality tickets than they would otherwise purchase. The price effect is smaller than the quantity effect, but still economically meaningful. We see this magnitude of response in both a per-transaction and a per-user comparison. Interestingly, Back-end Fee users purchase a lower total ticket quantity, conditional on making at least one purchase. While this response margin (the total number of tickets purchased) is outside of the discrete choice model, it hints at the second effect of price obfuscation: a change in the marginal consumer. The marginal consumer under BF, attracted by seemingly lower prices, may be more price-sensitive and therefore in the market for fewer tickets compared to his UF counterpart.

Figure 1: Transaction Per Ticket Price Distribution

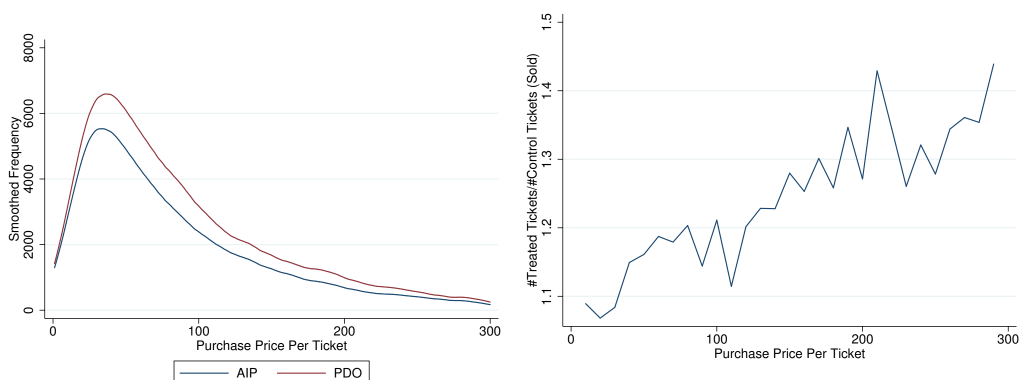

(a) Kernel Smoothed Histogram

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

We next consider the heterogeneity of response by price range. Theory suggests that Back-end Fee users should be more likely to purchase any ticket, but in particular, more likely to purchase an expensive ticket. Panel (a) in Figure 1a 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. There is a surprising consistency in response across all price points. But it is not completely uniform, as is clear in panel (b) of Figure 1b.

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.

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

5 Behavioral Mechanisms

5.1 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;[6] (3) the consumer reaches the checkout page where a final purchase decision is made; (4) the purchase confirmation page completes the process.[7] 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.

Table 5 shows the absolute and relative rate of UF and BF user arrivals at each step in this process (the signs are unchanged for the full sample computation, but we refrain from posting those figures to protect business-sensitive information). 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.

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

[7]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).

Table 5: Purchase Funnel Rates

	Relative to Initial Page				Relative to Prior Page			
	Upfront Fees	Back-end Fees	Difference	% Difference	Upfront Fees	Back-end Fees	Difference	% Difference
Event Page	100	100	–	–	–	–	–	–
Ticket Details	62.92	66.93	4.00	0.06	62.92	66.93	4.00	0.06
	(0.21)	(0.21)	(0.3)		(0.21)	(0.21)	(0.3)	
Start Checkout	13.72	21.45	7.73	0.56				
	(0.15)	(0.18)	(0.23)					
Review & Submit	10.04	15.66	5.61	0.56	15.96	23.40	7.43	0.47
	(0.13)	(0.16)	(0.21)		(0.2)	(0.23)	(0.3)	
Purchase	6.00	6.78	0.79	0.13	59.71	43.32	-16.39	-0.27
	(0.1)	(0.11)	(0.15)		(0.68)	(0.55)	(0.87)	

Table 6 examines the quality response at each step in the purchase funnel for a subset of events. The average price of tickets under 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 those with experience should not react to obfuscation.

To measure experience, we calculate the number of visits each cookie has made to StubHub.com. 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.[8] 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. Table 7 shows that the treatment effect is smaller for cookies with at least 20 site visits: the revenue effect is 15% compared to 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.

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

Table 6: Per Ticket Price Along the Purchase Funnel

	Back-end Fees	Upfront Fees	Difference	% Difference
Event Page	$1.00	$0.84	$0.16	19%
			(4.68)	
Ticket Details	$0.86	$0.78	$0.08	10%
			(7.60)	
Start Checkout	$0.77	$0.75	$0.02	3%
			(14.28)	
Review & Submit	$0.56	$0.52	$0.04	7%
			(5.34)	
Purchase	$0.42	$0.39	$0.03	7%
			(0.58)	

Notes: Average price of each touch point using that tickets price at time of touch. Standard errors in parentheses.

Table 7: Salience by User Experience

	% Difference		
	New User	Low Experience	High Experience
Cookie 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%

5.2 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. Table 8 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 9 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. Back-end Fee users are twice as likely to view 3 or more listings than

Table 8: Average Number of Tickets Viewed

# of Tickets	PDO	AIP
1	75%	84%
2	16%	12%
3	5%	3%
4	2%	1%
5+	2%	1%

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

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 Ellison (2009), who find that firms endogenously create such frictions to soften price competition.

Table 9: 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%)

Standard errors in parentheses.

5.3 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 9 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.

5.4 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 2 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.

Figure 2: Google Queries for Competing Ticket Resale Platforms

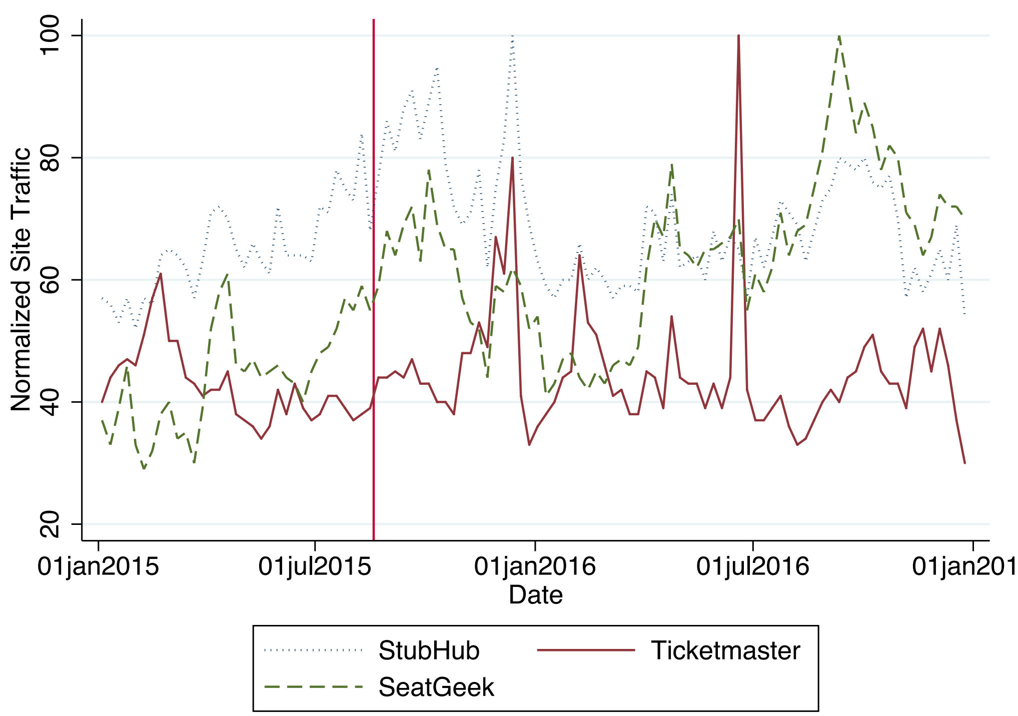

6 A second experiment: randomization at the event-level

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 10: 2012 Experiment Resultsl

	Upfront vs Back-end Fees % Difference
Purchase Probability	-0.13
	(0.07)
Percentile of Choice Set Selected	-0.12
	(0.06)

Notes: Standard errors clustered at the event level.

The results from the 2012 MLS experiment, displayed in table 10 , 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,

Figure 3: Percentile of Choice Set Purchased in the 2012 Experiment

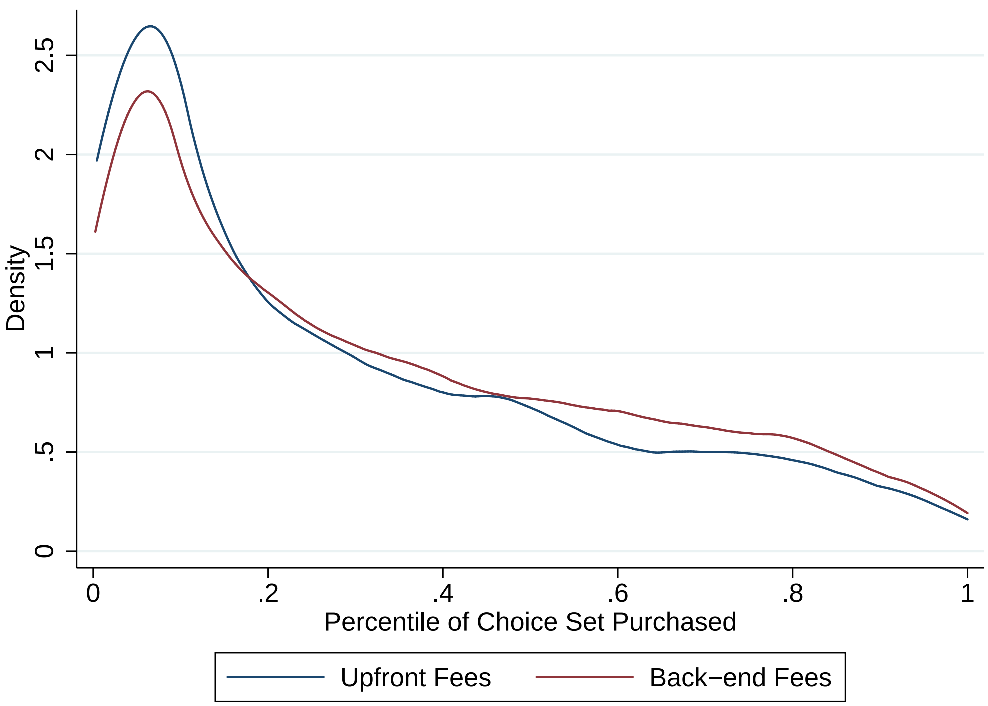

consumers buy from a 12% lower quantiles for UF compared to BF games. Figure 3 shows the full distribution of purchase quantiles for BF and UF matches.

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

7 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 slightly smaller, but still meaningful, price effect. 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.

Blake, Thomas, 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.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *The American Economic Review*, September 2009, *99* (4), 1145–1177.

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.