PRICE SALIENCE AND PRODUCT CHOICE

TOM BLAKE, DOMINIC COEY, KANE SWEENEY, AND STEVE TADELIS

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 comparable in size to the effect on quantity. We estimate a structural behavioral model of product choice that allows us to quantify how much consumers misperceive true prices under price shrouding, and to evaluate the welfare loss from quality and quantity distortions due to price misperceptions.

Date: Jan 20, 2015.
Very preliminary, comments especially welcome.
Blake: eBay Research Labs; thblake@ebay.com.
Coey: eBay Research Labs; dcoey@ebay.com.
Sweeney: eBay Research Labs; kane.sweeney@ebay.com.
Tadelis: eBay Research Labs; stadelis@ebay.com.

1. Introduction

Textbook economic models of consumer choice assume that economic agents are fully aware of the final prices of all goods, inclusive of fees and taxes, implying that consumption decisions are based on the true final prices of all goods. As a consequence, any change in the final prices of goods, whether due to a change in the base price, the fees, or the tax rates, will result in the same change in the choice behavior of consumers. Recent influential work, however, has challenged this assumption (Chetty, Looney, and Kroft (2009), Finkelstein (2009)), convincingly showing that the extent to which taxes or fees are salient will have profound effects on whether consumers purchase certain goods and services. In this paper we use a large scale field experiment to show that the effect of price salience on *what* consumers purchase can be just as important as *whether* they purchase.

Consider the example of a percentage fee levied on all the goods in a consumer's choice set. The extent to which the final price is salient affects the consumer along two margins. First, increasing the salience of the fee makes all goods appear more expensive, resulting in a higher probability that the consumer chooses not to buy any good. Second, because a change in a percentage fee has a stronger effect on more expensive goods, it will change the perceived price-quality tradeoff for the goods in the consumers choice set, and as a result may encourage the consumer to substitute between more and less expensive 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 of these margins.[1]

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 all listed upfront—then consumers will not only be less likely to purchase a good, but conditional on purchasing, they will purchase lower quality goods.

[1]In the working paper of Chetty, Looney, and Kroft (2007), they note that the revenue effect is bigger than the quantity effect, which is potentially due to consumers switching to lower priced items. Unfortunately, their data is insufficient to investigate that possibility further.

We take these predictions to data generated from a large-scale field experiment conducted by a leading online secondary ticket marketplace for sports events. On the marketplace website, consumers historically first saw the list-price of tickets for sale in a certain stadium for a certain game, and would only see ticket fees, and hence final prices, after selecting a seat and proceeding to the checkout page.

In 2012, the firm randomly selected 33 games from the Major League Soccer season to be in the "All-in-Pricing" (AIP) treatment, in which consumers see the all-inclusive prices up front, and these prices remain the same on the checkout page. Another 66 games were left in the control group where fees were tacked onto the price only at checkout. This exogenous variation in fee salience, together with rich data on the actual choice sets that consumers faced, allows us to infer the effect of salience on product choice. Our empirical results indicate that the models predictions are borne out in the data. Both the quality and quantity distortions from price shrouding are present, and more interestingly, these two effects have a comparable impact on revenue raised from fees.

To assess the degree of price misperception, and to evaluate the welfare cost of price shrouding, we proceed to estimate a structural model, which follows naturally from our simple theoretical model, of product choice in the presence of behavioral biases that arise from differing levels of price salience. Namely, we allow for changes in fees to have a different effect on product choices than changes in base prices, where the difference is attributable to misperception. Similar to Chetty, Looney, and Kroft (2009), our welfare analysis proceeds by assuming that under AIP, consumers are perfectly informed and capable of selecting their most preferred ticked.[2] This approach relies on two assumptions: fees affect welfare only through their effects on the goods chosen by consumers; and consumption choices, when prices are perfectly salient in the AIP treatment, are optimal. Using the estimates of the choice model from the AIP case, we can evaluate the welfare of the choices made in both the AIP and non-AIP cases.

[2]Like Chetty, Looney, and Kroft (2009), this approach does not rely on a specific positive model of behavior, as described in Bernheim and Rangel (2009).

The next section reviews related literature. Section 3 contains the theoretical model and derives empirical implications. Section 4 discusses the experiment run at a secondary ticket market platform as well as the data used in the analysis. Section 5 provides some preliminary, non-parametric results. To disentangle the quantity from quality effects and perform welfare analysis, section 6 introduces the structural behavioral model; i.e. a typical model augmented with "behavioral" parameters. Section 7 contains the results of the model estimation, and section 8 concludes.

2. Related Literature

2.1. **Empirical Work on Tax/Fee Salience.**

- Chetty, Looney, and Kroft (2009)
- Chetty (2009)
- Chetty, Looney, and Kroft (2007)
- Finkelstein (2009)
- Blumkin, Ruffle, and Ganun (2012)

2.2. **Behavioral IO.**

- Ellison (2006)
- Grubb and Osborne (2015)
- Handel (2013)

2.3. **Structural Behavioral Economics.**

- DellaVigna, List, and Malmendier (2012)
- Conlin, O'Donoghue, and Vogelsang (2007)
- Lacetera, Pope, and Sydnor (2012)

2.4. **Biased Beliefs.**

- Allcott (2013)
- Bollinger, Leslie, and Sorensen (2011)

3. Salience and Product Choice in Theory

In keeping with our empirical application, we use terms consistent with buying tickets on a platform to a sporting event. Specifically, a consumer arrives to the site looking for tickets to a game. She is then presented with the set of tickets available, J, with tickets varying by both price and quality.[3] The utility of a ticket $j \in J$ to consumer i is:

$$v_{ij} = \delta_j - \alpha p_j,$$

where α measures the consumer's willingness to trade off quality for money. For convenience, let $0 \in J$ denote the outside option, with $\delta_0 = p_0 = 0$. For clarity the traditional error terms are suppressed in this toy model, but included in the empirical specification below.

In a standard discrete choice model, the consumer would select the option with the highest utility:

$$\max_{j \in J} \delta_j - \alpha p_j \tag{1}$$

Let $j^* \in J$ denote the optimal choice, keeping in mind that the consumer may select $j = 0$; i.e. not purchase a ticket.

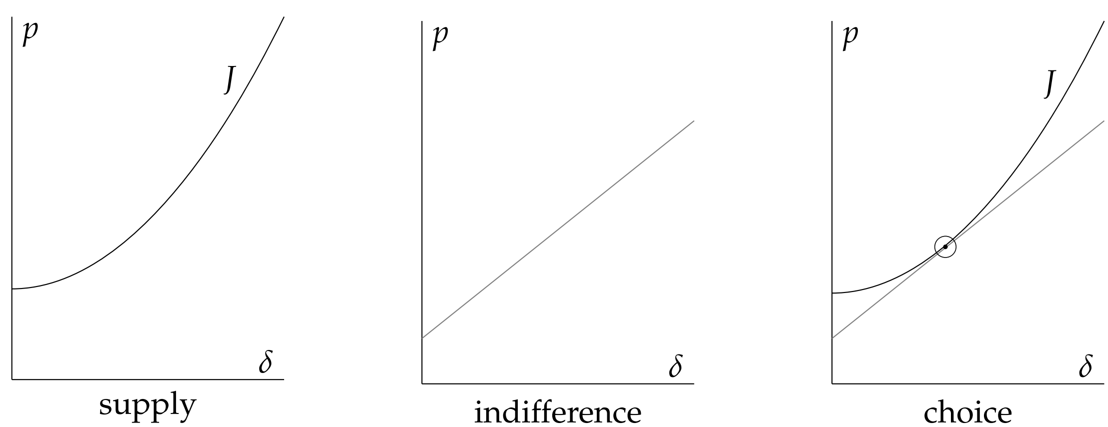

The panel on the left displays the supply of tickets available to a consumer, while the center panel illustrates an indifference curve. The consumer then chooses the

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

ticket from the supplied set on her highest indifference curve, indicated by the tangency on the right. Note that the consumer's utility is linear in price and quality, which is without loss of generality. The set of available tickets can have any shape, but is drawn as convex increasing to foreshadow the empirical findings below.

To investigate how salience may affect consumer choice, we now depart from the standard model. Following the approach in Finkelstein (2009), we allow each ticket j to have both a price p_j as well as a *perceived* price $\tilde{p}_j$.

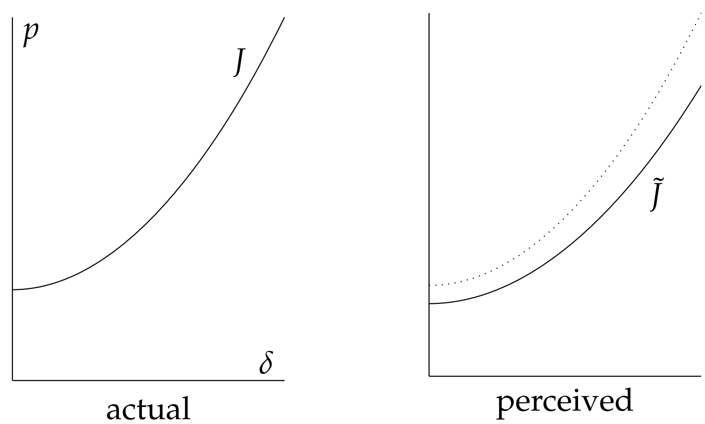

The price p_j is the full checkout price and is observable to the econometrician, while the perceived price is unobserved. However, to connect the theory to observables, we assume that the consumer maximizes her expected utility with respect to the observed quality and the perceived price. That is, she selects the $j \in J$ to solve the following optimization problem:

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

Let $\tilde{j}^*$ denote this observed choice. Note that the "observed" demand $\tilde{j}^*$ maximizes the perceived utility, while the "optimal" demand j^* maximizes the actual utility.

Motivated by the actual experiment, we now present two testable implications which relate perceived prices to empirical predictions. First, when the perceived prices are higher than the actual prices, we expect to see higher rate of non-purchase than if consumers percieve the prices correctly. This is the chief empirical finding in the existing literature: when fee-inclusive prices are made salient, total quantity purchased declines (Chetty, Looney, and Kroft (2009); Finkelstein (2009)).

Observation 1. *If perceived prices are lower than actual prices, "observed" demand is higher than "actual" demand; i.e. if*

$$\tilde{p} \leq p \quad \& \quad \tilde{j}^* = 0 \qquad \Rightarrow \qquad j^* = 0 \tag{2}$$

That is, if the perceived prices are lower than the actual prices ($\tilde{p} \leq p$)[4] and the consumer does not buy a ticket under the perceived prices ($\tilde{j}^* = 0$) then she would continue to not buy under the actual, higher prices ($j^* = 0$).

However, if the difference between the perceived prices and the actual prices is *increasing* in the actual prices (as one might expect with with percentage fees or taxes), then customers also buy lower quality items than they would under correct perceptions.

Observation 2. *If $\tilde{p} \leq p$ and $\tilde{p} - p$ is increasing in p, "observed" demand is lower quality than "actual" demand, conditional upon purchase:*

$$\tilde{j}^*, j^* \neq 0 \qquad \Rightarrow \qquad \delta_{\tilde{j}^*} \geq \delta_{j^*} \tag{3}$$

In the illustration below, we see the impact of the misperception: under the misperceived costs, the consumer buys a higher quality (and higher priced) ticket than she would under correctly perceived prices.

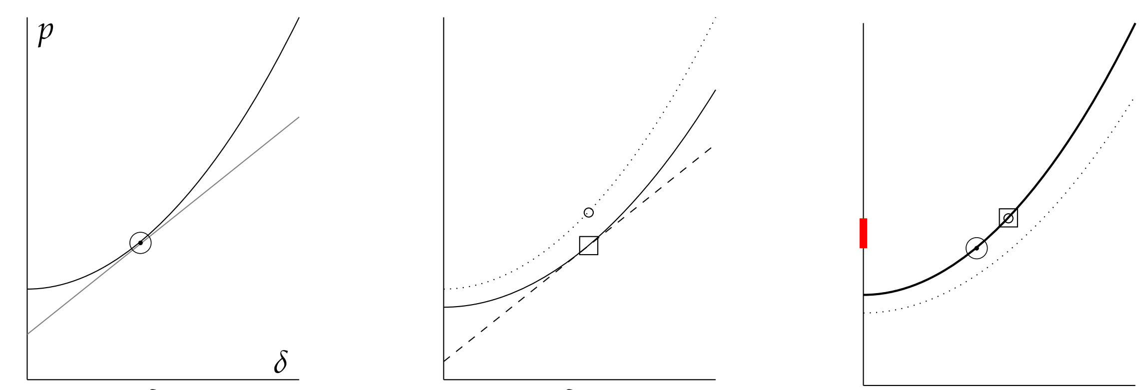

$\max \delta_j - \alpha p_j$ $\qquad$ $\max \delta_j - \alpha \tilde{p}_j$ $\qquad$ "observed" diff

[4]We write "$\tilde{p} \leq p$" instead of the more cumbersome "$\tilde{p}_j \leq p_j \, \forall j \in J$".

In the next two sections, we provide reduced-form evidence which is consistent with both Observations 1 and 2. In particular, we show that making fees more salient lowers both the quantity and quality of tickets sold.

4. Data and Experiment Overview

Prior to 2012, the platform added its buy fee at the checkout page, much like taxes added at the register of a grocery store.[5] By 2013, all fees were shown up front to consumers, so that the initial price that the consumer sees is the final checkout price. The platform refers to this as *All-in Pricing* (AIP). In 2012, an experiment was run during the MLS regular season. The season was split into treatment and control, with the treatment being AIP and control being the then-status quo; i.e. no AIP. 33 games were randomly assigned to have the AIP experience, while 66 games were designated as control.[6] In all games, the fee to the buyers was 10% of the ticket price. In the AIP games, conspicuous onsite announcements informed consumers that the prices they saw were the prices inclusive of all charges and fees. We consequently assume that the perceived price in the AIP treatment is equal to the actual price.[7]

[5] A similar timing of fee announcement is currently in use at other ticket platforms, including TicketMaster.

[6] The other MLS games were used for different aspects of the experiment.

[7] In follow-up work, we investigate this assumption using detailed browsing data.

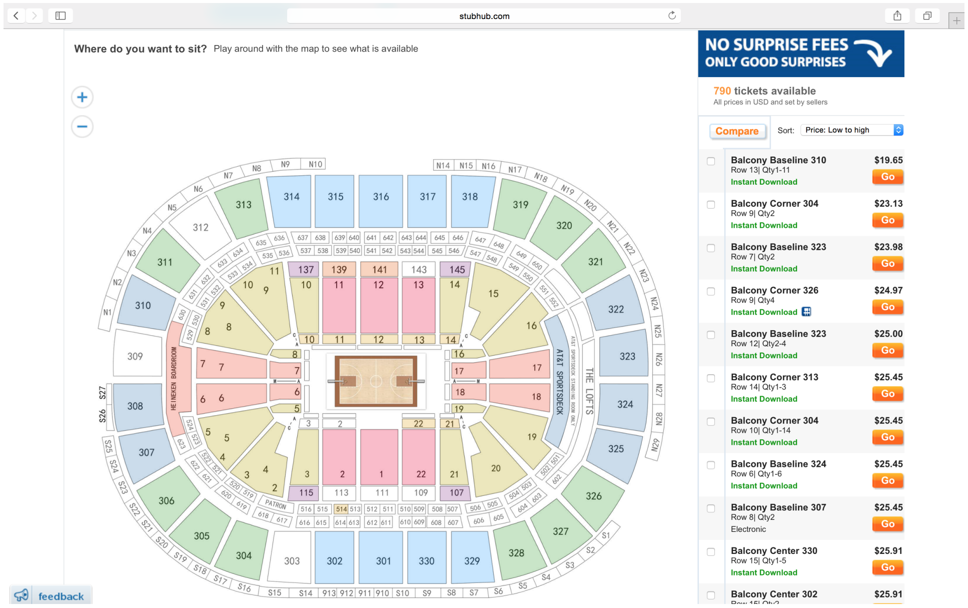

[SUMMARY STATS TO BE INCLUDED]

As a check on the randomization, below is a kernel density of consumers arriving to the site across AIP and non-AIP. The central message is that these densities are essentially identical; i.e. the arrival rates across treatment and control are similar. The green line has slightly higher variance, since it corresponds to the 33 games of AIP (relative to 66 games of non-AIP). Since the arrival rate of consumers to the site should not be affected by the experiment, it is reassuring that this is reflected in the data. The vast majority of consumer arrivals occur in the weeks prior to the event.

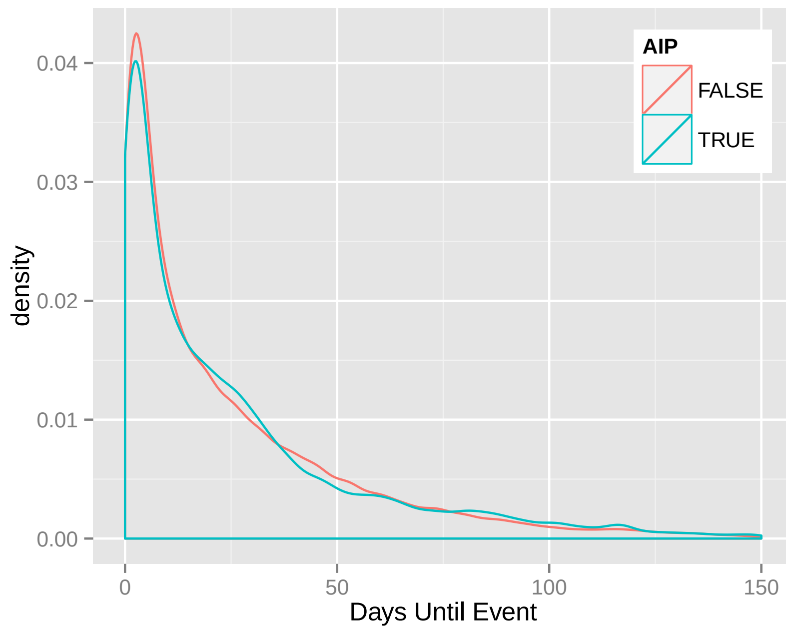

5. Experimental Results

If the hypotheses of Theorems 1 and 2 hold, then AIP would (i) encourage some consumers to switch from buying something to buying nothing and (ii) encourage some consumers to switch from higher to lower quality tickets. We now present the effects on those quantity and quality margins, as well as an estimate of the total effect.

5.1. **Quantity Effect of Salience.** We begin by examining the quantity impact, and the following graph shows two relationships. The bottom two curves represent the density of the arrival rate of people to the site, as described above.

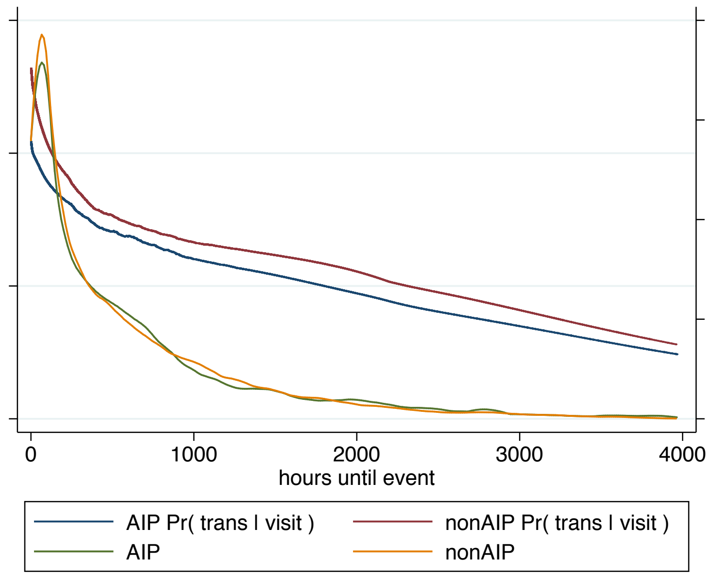

The upper two arcs (red and blue) represent the true quantities of interest: the probability that a consumer purchases a ticket, conditional upon arriving at the site. More specifically, these represent the probability that a consumer purchases a ticket for a game, once she has gotten to that game's ticket listings.

Similar to the arrival rate, consumers are significantly more likely to purchase as the event approaches. Averaging over this time window, consumers in AIP treatment are roughly 13% less likely to purchase a ticket during a visit in the AIP treatment, which is fairly consistent over time. This captures the quantity effect, and is roughly in line with results in the prior literature. Below we simply regress the probability of transaction on treatment, clustering the standard errors at the event level.

TABLE 1. Outcome: Purchase probability

Variable	Coefficient	(se)	(clustered se)
Intercept	0.078	(0.01)	(0.05)
AIP	-0.010	(0.01)	(0.05)

This graphic *does not* imply that the total revenue effect is 13%. The full effect could be more or less than 13%, depending on how the quality composition changes due to the treatment. We provide evidence of the importance of quality next.

5.2. **Quality Effect of Salience.** Relative to quality, the quantity effect is rather straightforward to assess: we can simply compare averages of the number of tickets sold across treatment and control. Assessing the impact on quality is more subtle, which is partially due to difficulties in measuring quality; i.e. if the true underlying quality were printed on the ticket, we would simply compare averages as we do with quantity.[8] To see provide some evidence that quality is important, we ask a related question:

> What percentile of her choice set did each consumer purchase, in terms of ticket price?

If showing consumers the 10% fee upfront encourages the difference between perceived prices and actual prices to be increasing in quality (Theorem 2), we should expect to see a stochastic shift in the purchasing behavior from higher to lower quality tickets *within their own choice set*. Since price is positively associated with quality, we should expect a shift toward lower priced tickets, where price is measured in both treatments as the total checkout price. We present results on the percentile chosen because it is a simple measure which preserves monotonicity in price, and thus represents an imperfect (but useful) measure in the shift in quality. The main graphic is below.

Each dot in the top panel presents a single transaction; the x-axis is unordered and the y-axis indicates the percentile of her choice set that customer purchased. For example, along the bottom of the panel, each consumer is purchasing the cheapest ticket available. At the top, the consumer is buying the most expensive. This top panel is meant only as an aid to explaining the bottom panel; i.e. the bottom panel is a simple resorting of the top panel, which produces an inverse CDF.

[8]An alternative would be to simply compare average revenues, but this is infeasible due to the large revenue at the game level. More on this is below.

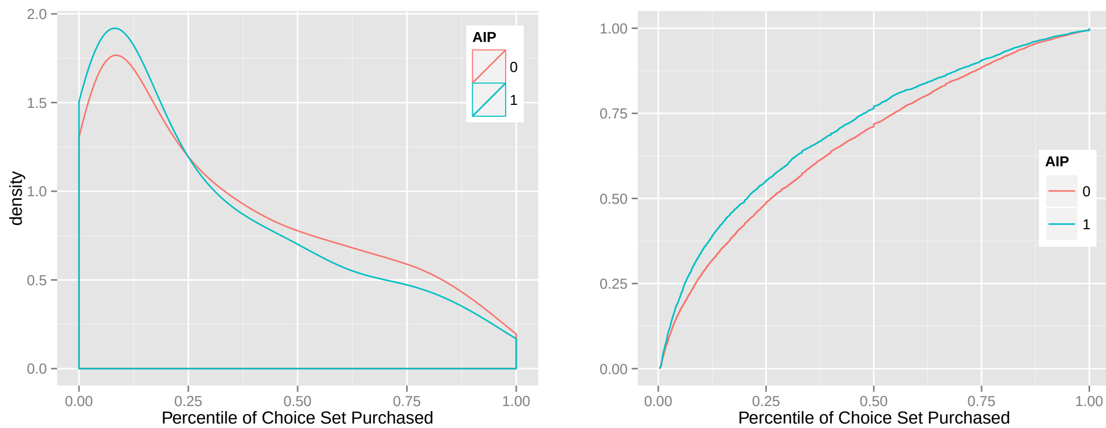

Consider the bottom panel with $x = 0.50$. Under AIP, this median consumer is purchasing the 26% most expensive ticket. The median non-AIP consumer is purchasing the 20% most expensive ticket. Recall that these are all *conditional upon purchase*, and reasoning about selected samples can be tricky. We look into this in the next section.

TABLE 2. Outcome: Percentile purchased (conditional on purchase)

Variable	**Coefficient**	**(se)**	**(clustered se)**
Intercept	0.336	(0.003)	(0.010)
AIP	-0.040	(0.005)	(0.019)

5.3. **Total Effect of Salience.** The most straightforward way to measure the total (revenue) effect of salience would be to simply compare the average revenue in the treatment games to the average revenue in the control games. This straightforward approach is unfortunately unworkable, since total game revenue varies substantially from game to game. The problem is essentially the same as the difficulty in measuring (revenue) returns to advertising: when aggregated to the game level, the signal is vastly smaller than the noise. However, from the previous two subsections, it should be clear that there is, in fact, a sizable effect from the treatment. The total effect cannot be stitched together from the two margin effects above (quantity and quality).

In the next section, we provide a more formal framework for estimating the total, quality and quantity effects. In this subsection, we impose a simple assumption which is unlikely to hold, strictly speaking, but will permit us to generate a useful benchmark. Specifically,

[RESULTS TO BE INCLUDED]

6. MODELING SALIENCE AND PRODUCT CHOICE EMPIRCALLY

The empirical results of Section 5 show that AIP affects both the likelihood of buying a ticket and the kind of ticket bought, but they do not tell us what is the welfare gain to consumers from price salience. To address this question we estimate a structural model of consumer choice, using the experimental variation in the AIP treatment. The key assumption for our welfare analysis is that under AIP, consumers are rational and fully informed, and so choose the ticket which maximizes their actual utility. We can use the estimates of the choice model obtained under AIP to evaluate the consumer welfare associated with any other set of choices. In particular, we can evaluate the consumer welfare of choices made when AIP is not in effect, and compare it to the welfare of choices made under AIP.

For game g, under either AIP or non-AIP, consumer i has an actual utility for ticket j of:

$$u_{ijg} = \delta_g + \delta_j + \alpha p_j + \varepsilon_{ijg}. \tag{4}$$

where δ_j is a function of section, row and delivery method, δ_g is a game-level fixed effect, and p_j is the actual total checkout price. The δ parameters capture the quality of the ticket, and the parameter α measures how consumers are willing to trade off price for quality.

In the experiment, consumers i is assumed to maximize her perceived utility $\tilde{u}$. Under AIP, consumers see the actual price upfront, so that $\tilde{u}_{ijg} = u_{ijg}$. If game g is non-AIP, then choices are determined by maximizing:

$$\tilde{u}_{ijg} = \delta_g + \delta_j + \alpha\left[\gamma + (1+\theta)p_j\right] + \varepsilon_{ijg}. \tag{5}$$

Relating back to an earlier theoretical model, under non-AIP, a consumer correctly observes ticket j's quality to be $(\delta_g + \delta_j)$ but she perceives the price to be

$$\tilde{p}_j = \gamma + (1+\theta)p_j.$$

In fact, we write (5) in its form to emphasize the same quality-price geometry from the earlier model in (δ, p) space. In the illustration below, γ and θ are both negative.

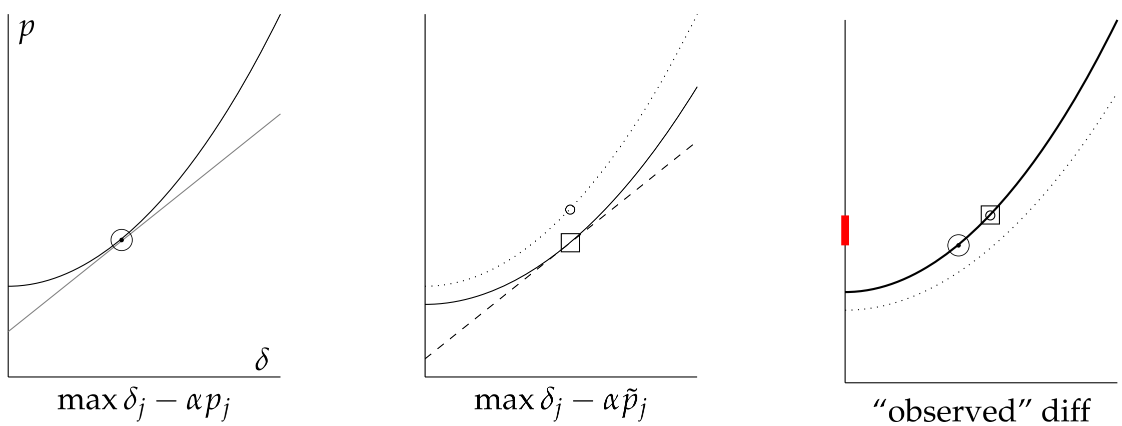

Thus, in general, choices are determined by maximizing:

$$\tilde{u}_{ijg} = \delta_g + \delta_j + \alpha\left[\gamma \mathbb{1}_g^{AIP} + (1+\theta\mathbb{1}_g^{AIP})p_j\right] + \varepsilon_{ijg}. \tag{6}$$

where $\mathbb{1}_g^{AIP}$ be an indicator for whether game g is AIP. Note that an increase in γ reduces the probability that any individual buys a ticket, but – conditional upon purchase – it does not affect her choice. Roughly speaking, γ is identified by the differential probability of purchase across the two treatments. The other behavioral parameter, θ, acts to tilt the perceived supply function, which impacts both the quantity and quality decisions. For a fixed γ, θ is identified by the difference in quality choices across treatments.

Again, the outside option of not buying a ticket is normalized to zero, and consumer i chooses the j maximizing $\tilde{u}_{ijg}$, as long as $\max_j \tilde{u}_{ijg}$ exceeds zero. If $\gamma = \theta = 0$, then consumers make the same choices with and without AIP.

6.1. **Welfare.** In AIP, we assume consumers are fully informed and capable of selecting the ticket that maximizes their utility. In non-AIP, we allow for the possibility that the price shrouding may induce consumers to choose tickets which do not maximize their utility. Thus, while (6) is valid as a positive description of consumer behavior under both non-AIP and AIP, only equation (4) captures the actual utility consumers receive from their choices.

Given the true utility function (equation (4)) and the observed non-AIP choices, the consumer welfare gain in switching a non-AIP game g to AIP can easily be computed. If

$$j^* = \arg\max_j u_{ijg} \qquad \text{and} \qquad j' = \arg\max_j \tilde{u}_{ijg} \tag{7}$$

then the welfare gain from switching to the more transparency AIP is

$$\frac{1}{\alpha}\left(u_{ij^*g} - u_{ij'g}\right). \tag{8}$$

Note that j' maximizes the perceived utility but is evaluated according to the true utility.

6.2. **Estimation.** We recover the parameters of equation (6) in two steps. We first define

$$\eta_g = \delta_g + \alpha\gamma\mathbb{1}_g^{AIP} \tag{9}$$

and rewrite (6) as:

$$\tilde{u}_{ijg} = \eta_g + \delta_j + \alpha\left[(1 + \theta\mathbb{1}_g^{AIP})p_j\right] + \varepsilon_{ijg}. \tag{10}$$

Putting in game fixed effects, we estimate $(\eta, \delta, \alpha, \theta)$ of equation (**??**) using standard discrete choice techniques; e.g. maximum likelihood of the choice probabilities.

In a second step, we regress the estimated $\{\eta_g\}$ onto the AIP indicator:

$$\eta_g = \beta_0 + \beta_1 \mathbb{1}_g^{AIP} + \omega_g \tag{11}$$

where $\mathbb{E}[\omega_g] = 0$. We finally recover γ from

$$\gamma = \frac{\beta_1}{\alpha} \tag{12}$$

Equations (10) and (11) generate our two estimating equations. Estimation can proceed by a two step procedure, in which we first estimate the discrete choice logit model in equation (10), and then regress the estimated δ_g's on the $\mathbb{1}_g^{AIP}$ indicator in a second stage. Alternatively, we can form moment conditions for both equations and estimate all parameters simultaneously, in a GMM procedure.

We do not anticipate endogeneity being a key issue in the estimation, either for estimating the price sensitivity coefficient or the effect of AIP. We control for essentially all the information that buyers observe at the time of the purchase decision, which mitigates the concern that prices might be correlated with product characteristics unobservable to the econometrican, but observable to the buyer. The exogenous variation in prices conditional on these observables comes from heterogeneity in sellers' eagerness to sell.[9] Further, the AIP treatment is randomly assigned across games, and therefore uncorrelated with the errors ε_{ijg} and ε_g.

7. Conclusion

We have introduced a choice model with both product quality variation and sensitivity to salience. The main result of the simple theoretical model is that consumers should tend to buy lower quality items when a change in salience renders perceived prices relatively higher for higher quality items, which is plausible given that fees are often multiplicative or percentages. We then show that this conjecture is supported by data in an experiment run by a secondary ticket market platform.

[9]We do not include seller characteristics: the ticket platform guarantees all transactions, and buyers do not know the identity of sellers.

One major caveat to this analysis is that these results should be interpreted as measurements of short run effects. A platform is arguably willing to institute a policy which has short run negative effects if the longer run effects are positive. In the current setting, it is plausible (or even likely) that educating consumers about all-in pricing may require more time than just the experimental period of the 2012 MLS season. Education and experimentation are typically in tension because companies may not want to roll out a public campaign when a policy may only effect a small slice of consumer traffic. In ongoing work, we are analyzing the long run impact salience on the platform.

REFERENCES

ALLCOTT, H. (2013): "The welfare effects of misperceived product costs: Data and calibrations from the automobile market," *American Economic Journal: Economic Policy*, 5(3), 30–66.

BERNHEIM, B. D., AND A. RANGEL (2009): "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *The Quarterly Journal of Economics*, 124(1), 51–104.

BLUMKIN, T., B. J. RUFFLE, AND Y. GANUN (2012): "Are income and consumption taxes ever really equivalent? Evidence from a real-effort experiment with real goods," *European Economic Review*, 56(6), 1200–1219.

BOLLINGER, B., P. LESLIE, AND A. SORENSEN (2011): "Calorie Posting in Chain Restaurants," *American Economic Journal: Economic Policy*, 3(1), 91–128.

CHETTY, R. (2009): "The simple economics of salience and taxation," Discussion paper, National Bureau of Economic Research.

CHETTY, R., A. LOONEY, AND K. KROFT (2007): "Salience and taxation: Theory and evidence," Discussion paper, National Bureau of Economic Research.

——— (2009): "Salience and Taxation: Theory and Evidence," *The American Economic Review*, 99(4), 1145–1177.

CONLIN, M., T. O'DONOGHUE, AND T. J. VOGELSANG (2007): "Projection bias in catalog orders," *The American Economic Review*, pp. 1217–1249.

DELLAVIGNA, S., J. A. LIST, AND U. MALMENDIER (2012): "Testing for Altruism and Social Pressure in Charitable Giving," *Quarterly Journal of Economics*, 127(1), 1–56.

ELLISON, G. (2006): "Bounded rationality in industrial organization," *Econometric Society Monographs*, 42, 142.

FINKELSTEIN, A. (2009): "E-ztax: Tax Salience and Tax Rates*," *The Quarterly journal of economics*, 124(3), 969–1010.

GRUBB, M. D., AND M. OSBORNE (2015): "Cellular service demand: Biased beliefs, learning, and bill shock," *The American Economic Review*, 105(1), 234–271.

HANDEL, B. R. (2013): "Adverse selection and inertia in health insurance markets: When nudging hurts," *The American Economic Review*, 103(7), 2643–2682.

LACETERA, N., D. G. POPE, AND J. R. SYDNOR (2012): "Heuristic thinking and limited attention in the car market," *The American Economic Review*, 102(5), 2206–2236.

APPENDIX A. PROOFS