Irregular Demand: The Anti-Compromise Effect in the Field

Nicola Gennaioli, John A. List, Matthias Rodemeier, and Gregory Sun

November 2026

Abstract

Neoclassical demand obeys *regularity*: adding a good can only draw demand away from existing goods. In a natural field experiment with Lyft covering 7.2 million riders and 47 million choices, we randomize the presence and price of a premium ride. Contrary to regularity, the premium ride raises demand for the cheapest incumbent and for not buying. The effect grows with the premium price and with repeated exposure, and persists in sessions where the premium ride is not displayed. Rational inference and inattention cannot account for these facts. A model in which a surprisingly high price, cued on the app and retrieved from memory, draws attention to prices explains them.

JEL Codes: D01, D12, D9
Keywords: Revealed Preference, WARP, Natural Field Experiment

---

Affiliations: Gennaioli and Rodemeier: Bocconi University and IGIER. List: University of Chicago, Australian National University, and NBER. Sun: Olin Business School, Washington University.

1 Introduction

Deciding which products to offer is key for firms. Airlines carve economy into basic and premium tiers; smartphone makers add flagships; ride-hailing apps charge extra for a faster pickup. Such moves shape the menu consumers face. How do they affect choice? In neoclassical theory a new product allows a firm to better match some consumers' tastes, drawing demand from the firm's other products, rivals, and from not buying anything. Consumers gain, and so does the firm if the product is well chosen (Tirole, 1988). What a new product cannot do is to make any existing option more attractive. For deterministic utility this is the Weak Axiom of Revealed Preference (Samuelson, 1938): a consumer facing a larger menu either switches to the new good or keeps her previous choice. With random utility, it is the weaker regularity condition: adding an alternative lowers the choice probability of each incumbent option, the outside option included. These properties shape the effect of new goods on demand, profits and welfare.

We test them in a natural field experiment with the ride-sharing company Lyft, randomizing 7,242,942 riders across 47,616,543 choice occasions, none of whom knew they were participating in a study. Control riders chose between "Wait & Save," cheap but with a long pickup window, and "Standard," faster and more expensive. Treated riders saw the same two products at the same prices, plus a premium option, "Fast Pass," faster and more expensive still. We also randomized the Fast Pass price across three levels, always above Standard, holding its pickup time fixed. In the neoclassical model the vertically superior Fast Pass attracts riders with a high value of time, reducing the share choosing Standard and possibly also the shares choosing Wait & Save and not buying.

We detect a violation of this prediction taking three forms. First, adding Fast Pass raises the share of the *cheaper* Wait & Save option and of not buying (while decreasing the share of Standard). One interpretation is that Fast Pass raises riders' price elasticity,

moving them toward cheaper options. Second, the effect is stronger for higher Fast Pass prices, which therefore raise riders' price elasticity. Third, the effect also grows stronger over time; Experience with Fast Pass raises price elasticity. Magnitudes are meaningful. In a rider's first exposure, demand for Wait & Save rises by about 2 and 3 percent of its baseline share, while the probability of making no request rises by a more modest 0.3-0.5 percent; across all sessions, the Wait & Save effect is roughly five times larger. The effect of adding Fast Pass in a given pricing condition on the share of not buying equals that of a uniform 1.17 percent increase in every Lyft price; the effect of Fast Pass on the share of Wait & Save equals that of a 1.25 percent *decrease* in this option's price.

The violation of regularity shows that, contrary to standard screening logic, the premium tier did not move riders up the menu: it moved them down. It also differs from violations of WARP such as the compromise and decoy effects (Simonson, 1989; Huber et al., 1982), in which adding a premium product should favor more expensive incumbents; here it favors the cheaper one and not buying. What causes these effects? One rational account is information: the menu alters riders' inference about the firm or situation (Wernerfelt, 1995; Kamenica, 2008). These models, however, generate the compromise effect, not its opposite. We examine other forms of menu-based inference and argue that they fail too. A chief reason is that rational riders would learn that Fast Pass and its price are random, and hence uninformative, so the effect should fade over time. Instead, it grows. Another account is rational inattention: consumers optimally choose what to learn about their options, so a new good shapes valuation of incumbent goods (Matějka and McKay, 2015). Here Fast Pass should matter especially if it is worth learning about, so the violation of regularity should shrink when Fast Pass is more pricey and hence chosen less often. In the data the effect is largest at the highest price, where Fast Pass is chosen by only one percent of riders.

We show that the data can be explained by a model of salience where, following Bordalo et al. (2020, 2026), current cues and past experiences prompt a rider to represent choice as a "large price variation" situation, increasing her bottom-up attention to prices (Kőszegi and Szeidl, 2013; Bordalo et al., 2013). The model produces our three findings. First, upon seeing the expensive Fast Pass, the rider recalls expensive rides from the past, which increase her perceived price variation. She thus pays more attention to price, increasing her price elasticity. Second, the effect is stronger when Fast Pass is more expensive, increasing price elasticity compared to a lower Fast Pass price. Third, experience with Fast Pass fills memory with high prices, causing perceived price variation and price elasticity to grow over time.

We also exploit the fact that in some treated sessions Fast Pass was not available due to limited supply. Even in these sessions, riders treated with high Fast Pass prices exhibit a higher price elasticity than those treated with low ones. Salience can produce this fact: the former riders recall higher price variation and thus exhibit higher price elasticity than the latter riders. Consistent with memory, this effect is however weaker than the effect detected when Fast Pass is available.

We contribute to work on violations of revealed preference (e.g., Simonson and Tversky, 1992; Tversky and Simonson, 1993; Huber et al., 1982; Herne, 1999; Kroll and Vogt, 2012; Crosetto and Gaudeul, 2016; Andreoni and Miller, 2002; Somerville, 2022). Much of it relies on laboratory experiments, leaving open whether violations survive in real markets, where experience and self-selection can attenuate anomalies (e.g., List, 2003; Enke et al., 2023; Rodemeier, 2025; Miller et al., 2026). Two natural field experiments add a premium option in an online bookstore (Gu et al., 2018) and recommend it on a travel aggregator (Rafai et al., 2022); the first finds a compromise effect, the second no effect.[1] We present a large-scale natural field experiment that randomizes the presence of

[1] Cai (2026) offers different insurance menus to Chinese farmers and also finds a compromise effect:

a premium good and its price, and that exposes the same rider to the new menu repeatedly over time. We find an anti-compromise effect: the premium good raises price elasticity, sorting consumers into the cheaper option and into not buying, the more so the higher its price and the longer the exposure. Our findings relate to Gennaioli et al. (2026). Their field violation of revealed preference can also be explained by their treatment (nonlinear vs linear pricing) increasing price salience and elasticity.

2 Experiment and Results

2.1 Design and Sample

The experiment ran from November 5, 2020 to January 20, 2021. Randomization was at the rider level. In Control, riders are offered "Wait & Save" (henceforth, "W&S"), which has a low price but requires riders to accept a longer pickup window, and the normal option "Standard" which has a shorter, precisely stated pickup time at a higher price. Riders could choose not to buy anything. In the treatment group W&S and Standard are unchanged but we introduce a premium product called "Fast Pass," which offered a shorter pickup time than Standard, but with a price markup. The expanded menu thus contained three ordered price–time alternatives: W&S, Standard, and Fast Pass. We also randomized the Fast Pass price into Low, Medium and High.[2] Any user received the same markup in all sessions. Figure 2 shows the visual layout riders saw in the Lyft app.

We observe 7,242,942 riders that open 47,616,543 sessions (a session is a discrete interaction where a consumer opened the Lyft app, input the origin and destination of

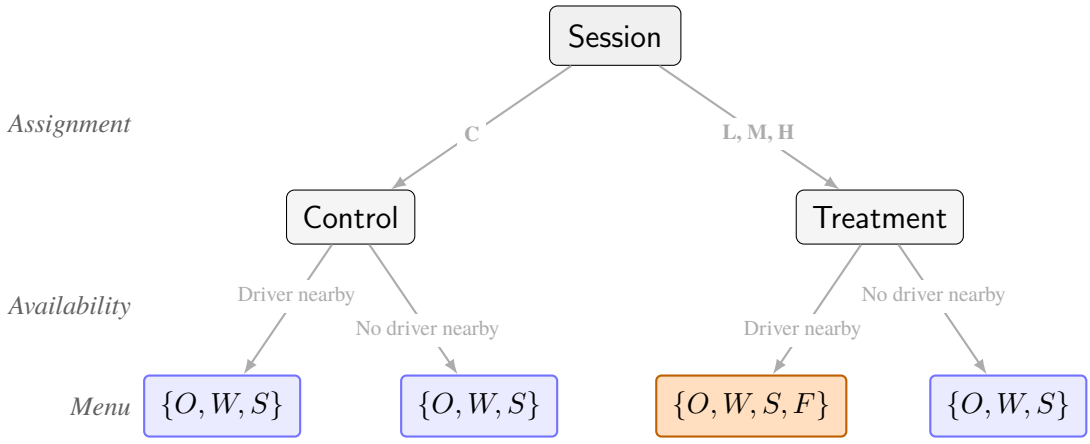
the premium option raises the probability of buying, the opposite of what we find. Her insurance is sold at organized village meetings rather than on a commercial platform. Work on choice overload (Bertrand et al., 2010; Long et al., 2025) varies the information displayed, not the menu consumers choose from.

[2]The Fast Pass markup over Standard prices varied across quintile of ride distance. The markup in the High Price treatment was roughly 4 times the markup in the Low Price treatment.

a ride, and received a price quote).[3] In the treatment group, Fast Pass could only be offered when a nearby driver was available. Fast Pass was available in approximately 33% of treated sessions, depending on the day of the week.

This experiment has two key features. First, the randomized Fast Pass and its price are orthogonal to types of rider, time conditions, including those determining the availability of Fast Pass, see Figure 1. This implies that by comparing a treated and control rider we obtain a causal intent-to-treat effect of making Fast Pass available vs. not and at different prices. Second, the treatment is at the level of the individual rider and persistent, allowing to study a rider's repeated exposure to Fast Pass.

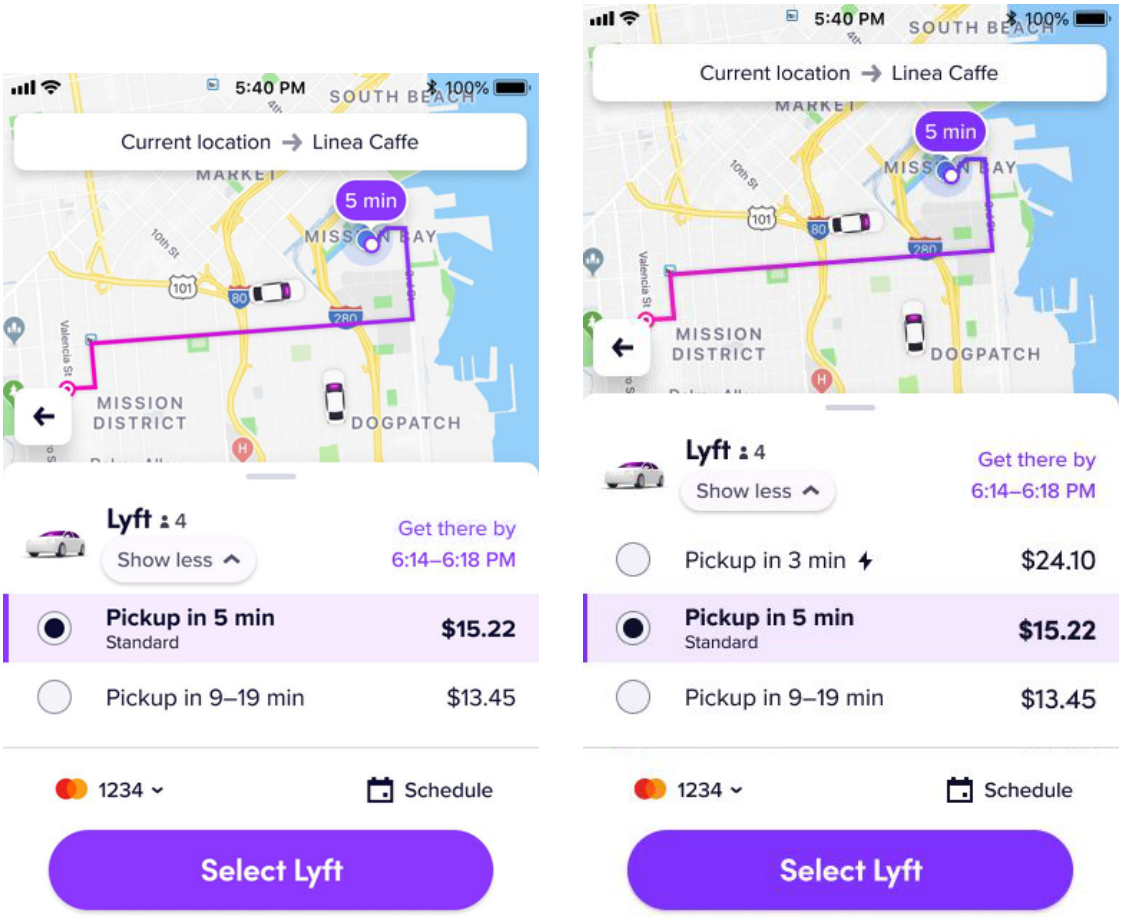
Figure 1: Experiment Design



Notes: This figure illustrates how experimental treatment and session characteristics jointly determine the menu available to the consumer. Experimental treatments are denoted by C, L, M, H, respectively denoting Control, Low Price, Medium Price, and High Price. The options O, W, S, F in the menu represent respectively the Outside Option, Wait & Save, Standard, and FastPass.

In an unreported analysis (available upon request) we show that randomization was successful. Our observables, namely lifetime rides and sessions, activity during the preceding 28 days, and pre-experimental request rates for Standard and Wait & Save, exhibit small and statistically insignificant differences across treatment and control groups.

[3] The control group contained 1,013,072 riders. The low-, medium-, and high-price Fast Pass treatments contained 2,463,821, 2,463,278, and 1,302,771 riders, respectively.



(a) Control **(b)** Treatment

Figure 2: Illustrative ride-option screens in the control group without the premium option (Panel A) and in the treatment group with the premium option showing a shorter pickup time of 3 minutes at a higher price of $24.10 (Panel B).

2.2 First Session Treatment Effects

The cleanest test of adding Fast Pass comes from the *first* session the consumer faces after randomization. Until this point, the consumer experience could not have been affected by treatment assignment. We run regressions

$$Y_i^j = \alpha^j + \sum_{p \in \{L,M,H\}} \beta_p^j \mathbf{1}\{Z_i = p\} + \varepsilon_i^j, \tag{1}$$

where $Y_i^j \in \{0,1\}$ indicates whether rider i chose option j = W,S,F,O in their first session, where j = O is the "Overall" demand for Lyft, so $1 - Y_i^{\text{O}}$ is the share of not buying. The treatment-indicator Z_i, indexes with $p \in \{L, M, H\}$ the low-, medium-, and high-price Fast Pass prices. The omitted category is Control.

Coefficient β_p^j captures the intent to treat effect of being assigned to a Fast Pass treatment. Regularity predicts lower demand for incumbent options $\beta_p^W, \beta_p^S \leq 0$, and $\beta_p^{\text{O}} \geq 0$. The effect should be smaller (β_p^j should be closer to zero) at higher Fast Pass prices.

Table 1 presents treatment coefficients for four outcomes: no request, Wait & Save, Standard, and Fast Pass.[4] Columns (1) and (2) show that, as expected, availability of Fast Pass increases demand for it and reduces demand for Standard, whose share falls significantly by 2.4 p.p., attenuating to 1.8, and 1.4 as we raise the Fast Pass price to medium or high. The corresponding Fast Pass demand is 2.5, 1.5, and 1.0 p.p.

In Column (3) being in medium and high Fast Pass price treatment increases demand for W&S by 0.2 and 0.3 p.p., respectively, with no effect in the low price treatment. On average, then, demand for W&S increases by 0.16 p.p. This increase in the demand for the cheaper incumbent option is a stark violation of regularity, in the opposite direction of the compromise effect.

[4]The total number of observations is smaller than the total number of randomized riders because a small fraction of riders never opened the app once throughout the entire experiment period.

Table 1: Request Rate Results: First Session

	Dependent variable:			
	Fast Pass (1)	Standard (2)	Wait & Save (3)	Overall (4)
Low FP Price	0.025***	-0.024***	-0.0002	-0.001*
	(0.0002)	(0.001)	(0.0004)	(0.001)
Med FP Price	0.015***	-0.018***	0.002***	-0.002***
	(0.0002)	(0.001)	(0.0004)	(0.001)
High FP Price	0.010***	-0.014***	0.003***	-0.002***
	(0.0002)	(0.001)	(0.0004)	(0.001)
Constant (No FP)		0.448***	0.107***	0.594***
		(0.001)	(0.0003)	(0.001)
Observations	6,587,607	6,587,607	6,587,607	6,587,607
R^2	0.004	0.0002	0.00001	0.00000
Adjusted R^2	0.004	0.0002	0.00001	0.00000
Residual std. error	0.124	0.495	0.310	0.491
F statistic	9,925.885***	545.902***	27.620***	6.268***

Standard errors, clustered by rider, are in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

In Column (4) Fast Pass reduces the probability of buying anything, or equivalently increases demand for the outside option, which also violates regularity. The effect also exhibits a price gradient: at medium or high Fast pass price the share of the outside option increases by 0.1 p.p. compared to the low-price treatment.

These effects map to a 1.5% increase of demand for W&S and a 0.41% increase in the probability of not buying compared to their 10.7% and 40.6% baseline. They materialize in the at most 33% of treated sessions with Fast Pass available, entailing a 4.5% increase in W&S and a 1.24% increase of not buying in these sessions.[5] They are equivalent to lowering the W&S price by 1.25% and to raising all Lyft prices by 0.89%, respectively.[6]

[5]Here we divide our observed intent-to-treat effect by the 33% probability of Fast Pass availability, averaging across all sessions. Ideally, we should divide by Fast Pass availability on the *first session* of each user, which we no longer have the micro-data to compute.

[6]Here we used ongoing price experimentation at Lyft, which gives a short-run demand elasticity of W&S to a 1% increase in own price equal to -3.6, and an elasticity of buying the outside option with

2.3 Effects on Subsequent Sessions

We can observe the evolution of the treatment effect at the rider level by looking at choice on sessions after the first one. To do so, we run regressions

$$Y_{is}^{j} = \alpha^{j} + \sum_{p \in \{L,M,H\}} \beta_{p}^{j} \mathbf{1}\{Z_i = p\} + \varepsilon_{is}^{j}, \tag{2}$$

where s is the session. Table 2 is the analog of Table 1. When looking at all sessions we continue to see violation of regularity that increase in Fast Pass price, but with magnitudes that are generally larger than in the first session. In particular, the increase in W&S demand in Column (3) is now between 4 and 5 times larger than in the first session.

Table 2: Request Rate Results: All Sessions

	Dependent variable:			
	Fast Pass (1)	Standard (2)	Wait & Save (3)	Overall (4)
Low FP Price	0.033***	-0.038***	0.006***	0.0002
	(0.0001)	(0.001)	(0.001)	(0.001)
Med FP Price	0.018***	-0.030***	0.011***	-0.002***
	(0.0001)	(0.001)	(0.001)	(0.001)
High FP Price	0.011***	-0.027***	0.013***	-0.003***
	(0.0001)	(0.001)	(0.001)	(0.001)
Constant (No FP)		0.489***	0.158***	0.676***
		(0.001)	(0.001)	(0.001)
Observations	47,616,543	47,616,543	47,616,543	47,616,543
R^2	0.007	0.001	0.0001	0.00001
Adjusted R^2	0.007	0.001	0.0001	0.00001
Residual std. error	0.137	0.498	0.372	0.469

Each column reports a session-level OLS regression on indicators for assignment to the three Fast Pass price treatments. The omitted category is assignment to the menu without Fast Pass. Standard errors, clustered by rider, are in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Identification here is less clean: seeing Fast Pass may change how often, and when,

respect to a 1% uniform price increase equal to 1.4.

a rider subsequently opens the app. Can the patterns in Table 2 arise from selection of rational riders? In one version of this story, treated riders who keep opening the app may have different tastes from those who stop. This cannot produce our results: news that Lyft offers a premium option should retain riders with a high value of time, which predicts more Standard and more purchases, not fewer.

Second, Fast Pass may deter a treated rider to open the app when she is in a hurry and expects to pay for speed, in which case Table 2 oversamples unhurried occasions, associated with a higher share of W&S and not buying. Two observations speak against this selection.[7] On the one hand, deterrence cannot occur with rational riders. A menu with Fast Pass weakly dominates one without it, since Standard and W&S are offered at the same prices, so seeing Fast Pass cannot deter a rational rider from subsequently opening the app. On the other hand, we see no deterrence. Our 6,229,870 treated riders opened 40,969,450 sessions and our 1,013,072 control riders 6,647,093, that is, 6.576 and 6.561 sessions per rider: a slightly higher 0.23 percent frequency of opening the app in treatment. Selection could then operate only if exposure raised app-opening on unhurried occasions. This is implausible: why should a low value of time rider be more interested in Lyft after seeing Fast Pass? Rational selection is unlikely to generate Table 2.

2.4 Effect of Fast Pass Availability

We finally ask: did prior exposure to Fast Pass change rider behavior even in sessions where Fast Pass is currently not available? One problem here is that availability depends on conditions (e.g. traffic) that may affect demand. Ideally, one would compare treated sessions in which Fast Pass was not available to Control sessions in which it would also have not been available. Unfortunately, storing this information at scale is very costly, so

[7] A third does too: if hurried occasions dropped out of the panel, demand for Fast Pass, which is bought on those occasions, should fall over time. Instead it grows, see Column (1) of Tables 1 and 2.

Lyft did not record it.

We deal with this problem by comparing treatment sessions in which Fast Pass was not available to treatment sessions with a low Fast Pass price in which Fast Pass was also not available. This strategy holds Fast Pass availability conditions constant and identifies the effect of the random Fast Pass price. For comparison we implement the strategy also for sessions where Fast Pass was available. Formally, we run regressions

$$Y_{is}^{j} = \alpha^{j} + \beta_{M}^{j}\mathbf{1}\{Z_i = M\} + \beta_{H}^{j}\mathbf{1}\{Z_i = H\} + \varepsilon_{is}^{j}. \tag{3}$$

separately on the sessions where Fast Pass was not available (resp. available). In these regressions, the constant α^j is the request rate in the omitted low-price Fast Pass group, β_M^j and β_H^j capture differences for the medium- and high-price groups.

Table 3: Request Rates by Fast Pass Availability

	Fast Pass not available			Fast Pass available			
	Standard	Wait & Save	Overall	Fast Pass	Standard	Wait & Save	Overall
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Medium FP price	-0.003***	0.003***	0.0001	-0.046***	0.029***	0.010***	-0.006***
	(0.001)	(0.0005)	(0.0005)	(0.0003)	(0.001)	(0.001)	(0.001)
High FP price	-0.005***	0.004***	-0.001**	-0.068***	0.043***	0.014***	-0.008***
	(0.001)	(0.001)	(0.001)	(0.0003)	(0.001)	(0.001)	(0.001)
Constant (low FP price)	0.488***	0.157***	0.672***	0.100***	0.375***	0.176***	0.683***
	(0.0004)	(0.0003)	(0.0003)	(0.0002)	(0.0004)	(0.0004)	(0.0004)
Observations	27,460,368	27,460,368	27,460,368	13,509,082	13,509,082	13,509,082	13,509,082
R^2	0.00001	0.00002	0.00000	0.012	0.001	0.0002	0.0001
Adjusted R^2	0.00001	0.00002	0.00000	0.012	0.001	0.0002	0.0001
Residual std. error	0.500	0.366	0.469	0.250	0.489	0.387	0.467

Notes: Each column reports a separate session-level panel regression with the low-price Fast Pass group as the omitted category. Standard errors are clustered at the rider-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

The left panel of Table 3 shows that past exposure to Fast Pass changes behavior even when Fast Pass is not on the menu. In columns (1)–(3), riders exposed to a medium and high Fast Pass price choose Standard less compared to those exposed to a low price, by

0.3 and 0.5 p.p. respectively. They also choose W&S more, by 0.3 and 0.4 p.p., and not buying more, by 0.1 p.p. Since Fast Pass is unavailable, the fall in Standard cannot reflect substitution: riders move down the menu as in Table 1, but now in response to a price that is absent from the screen and that, being random, carries no information. The effects are smaller than in the right panel, where the same price comparison is made with Fast Pass on the screen: about a third for W&S (1.0 and 1.4 percentage points there) and less than that for overall demand (0.6 and 0.8).[8] Experience with a striking, yet fully random, Fast Pass price seems to change the way riders value Standard and W&S.

2.5 Mechanisms

The observed failure of regularity is inconsistent with neoclassical stable tastes. It suggests that Fast Pass may increase riders' price elasticity, promoting the cheaper W&S and not buying. Three questions arise: 1) why is the effect stronger if Fast Pass is dearer? 2) Why does it get stronger with repeated exposure to Fast Pass? and 3) why should exposure to higher Fast Pass prices promote W&S and not buying even if Fast Pass is not available? We discuss theories of context effects and their consistency with our findings.

Decoy and Compromise effects Huber et al. (1982) document a "decoy effect" in the laboratory: choice of B increases relative to A when option C asymmetrically dominated by B is in the menu, a finding replicated by Herne (1999); Sonsino (2010); Kroll and Vogt (2012); Crosetto and Gaudeul (2016); Natenzon (2019). The idea is that the decoy C makes B look better by comparison. Simonson (1989) document a "compromise effect": adding an extreme option to the menu can increase demand for more intermediate options, which Bodner and Prelec (1994) and Kivetz et al. (2004) explain with loss-

[8] Standard is not comparable across panels: with Fast Pass displayed, a higher price also sends riders back from Fast Pass to Standard, which rises by 2.9 and 4.3 percentage points.

aversion. These effects cannot explain our findings because: i) Fast Pass is not dominated so decoy does not apply, and ii) the compromise effect would predict that adding Fast Pass should cause riders to move toward the middle option at Lyft, namely Standard, and reduce the outside option of not buying. Both effects are contrary to our findings.

Menu Based Inference Rational menu effects can arise from inference (Milgrom and Roberts, 1986; Wernerfelt, 1995; Kamenica, 2008). The firm knows more than riders — about their tastes, its own costs, etc. — and chooses the menu optimally, allowing rational riders to learn. This is unlikely to be first order in our setup: the app posts a price and a pickup time for every option (so little is left to infer) and riders are experienced, on average having 124 prior car-sharing sessions on Lyft alone. Suppose nonetheless that riders do infer, and consider signal structures varying across firms and time.

Time specific signal. The rider is uncertain about her current value of time, but she knows her population rank, i.e. whether she is a typically hurried or patient rider compared to others. Fast Pass may then signal that this is a peak time, leading the rider to revise her value of time up. Wernerfelt (1995) and Kamenica (2008) use this mechanism to explain the compromise effect. It implies that Fast Pass should cause riders to switch from W&S to Standard, which is counterfactual.

Alternatively, the rider may be uncertain about what posted time windows mean: the expected pickup time and its variability can differ over time. Fast Pass can signal (with its short window) that there are many drivers, leading to a faster arrival estimate. By this account, Standard and W&S should become more appealing. In the data, however, more riders choose not to buy. In principle, perceived strong supply may lead riders to consult other apps (e.g. Uber), but this could only foster not buying, not W&S.

Lyft specific signal. The Fast Pass price can signal that Lyft is expensive, prompting riders to check other apps. This signal could favor not buying over Fast Pass, but it is

redundant for Standard and W&S, whose price is shown in the app. Inconsistent with the data, then, demand should not be reallocated across these two products. Alternatively, upon seeing more options riders may infer that Lyft has not tailored a product for them, and withdraw (Kamenica, 2008; Kuksov and Villas-Boas, 2010). Again, this could explain a switch to the outside option but not from Standard to W&S.

Lyft-Time specific signal. Fast Pass may signal that Lyft is *currently* overcharging for speed. By reducing demand for speed, this would favor W&S, and not buying. The problem with this account is dynamic. The treatment was held fixed for eleven weeks, so a rider saw the same Fast Pass markup in all sessions. A rational rider would then learn that the markup was not related to current conditions, causing the treatment effect to wane. In reality, the effect grows stronger over time. In addition, and most directly, learning implies that the treatment effect should be absent or even reverse when Fast Pass is not shown. This is not true in the data: the demand for Wait & Save and for the outside option increases also in sessions where Fast Pass was never displayed.

Rational Inattention In Matějka and McKay (2015), when facing a choice menu consumer optimally choose what to learn about the current options. Here the presence of Fast Pass may prompt riders to learn more about their value of time, affecting choice among incumbent options, which could violate regularity. One problem of this mechanism is that it is stronger when Fast Pass is more attractive, so the effect should shrink as Fast Pass becomes more expensive and less chosen, which is not consistent with the data. A second problem is that this mechanism is absent when Fast Pass is not available, because the information strategy is menu-specific, inconsistent with Table 3.

A related account is efficient coding (Woodford, 2012): the rider perceives posted prices through a code adapted to experienced prices. Exposure to Fast Pass causes all Lyft prices to be perceived as higher, raising not buying. On the other hand, the extreme

Fast Pass price renders the W&S and Standard prices more similar, favoring Standard compared to W&S, a prediction also shared by relative thinking (Bushong et al., 2021). This is the opposite of what we find. A second problem is the role of experience. The model predicts no effect when Fast Pass appears for the first time, contrary to Table 1, and no difference between sessions with and without Fast Pass (the effect is much larger when Fast Pass is displayed).[9]

3 Salience Driven Attention

We show that a stylized model of salience based representations explains our findings based on the idea that the surprising Fast Pass price draws riders' attention toward prices bottom up, with an effect that grows and persists due to memory.

Following Bordalo et al. (2020, 2026), a rider's perception of the prices of options depends on current prices in the menu (the cues) and past experiences (memory). Specifically, the consumer chooses among the outside option 0, Standard S, W&S W, and Fast Pass F. In round $r = 0, 1, 2, \ldots$ a rider faces menu $\mathcal{M}_r \in \{\mathcal{C}, \mathcal{F}\}$, where $\mathcal{C} = \{0, \mathrm{W}, \mathrm{S}\}$ and $\mathcal{F} = \{0, \mathrm{W}, \mathrm{S}, \mathrm{F}\}$. The menu in Control and in the pre-experimental period $r = 0$ is $\mathcal{C}$. The treatment menu at $r \geq 1$, is $\mathcal{F}$ with i.i.d. probability $\phi \in (0, 1)$ and $\mathcal{C}$ otherwise.

Good j has speed s_j and price p_j, with $s_{\mathrm{F}} > s_{\mathrm{S}} > s_{\mathrm{W}} > s_0 = 0$ and $p_{\mathrm{F}} > p_{\mathrm{S}} > p_{\mathrm{W}} > p_0 = 0$. The value of good j to rider θ in round r is equal to:

$$u_j(\theta, r) = \theta \cdot s_j - \sigma_r \cdot p_j, \tag{4}$$

where θ, the rider's taste for speed, has an atomless cdf F and a density f with full

[9] In Caplin et al. (2011) inserting Fast Pass in the menu can push an incumbent out of the rider's consideration set, which under some assumptions could reproduce Table 1. But such effects require Fast Pass to be on the screen, so they cannot operate when Fast pass is not available, and they should fade as riders learn the new layout, whereas our effect grows.

support in $[0, \infty)$. The representation is set by $\sigma_r > 0$: the salience of price in round r. It affects valuation through the rider's weight on price vs speed. A rider with higher θ prefers faster options, but less so if prices are more salient, σ_r is higher.

Let $\pi_{\mathrm{W}} = p_{\mathrm{W}}/s_{\mathrm{W}}$, $\pi_{\mathrm{S}} = (p_{\mathrm{S}} - p_{\mathrm{W}})/(s_{\mathrm{S}} - s_{\mathrm{W}})$ and $\pi_{\mathrm{FP}} = (p_{\mathrm{FP}} - p_{\mathrm{S}})/(s_{\mathrm{FP}} - s_{\mathrm{S}})$ be the extra price per unit of speed of each good. We set $\pi_{\mathrm{W}} < \pi_{\mathrm{S}} < \pi_{\mathrm{FP}}$, so every good is chosen by some riders. Under salience σ, a rider picks 0 if $\theta < \sigma \cdot \pi_{\mathrm{W}}$, W if $\theta \in (\sigma \cdot \pi_{\mathrm{W}}, \sigma \cdot \pi_{\mathrm{S}})$, S if $\theta \in (\sigma \cdot \pi_{\mathrm{S}}, \sigma \cdot \pi_{\mathrm{FP}})$, and F if $\theta > \sigma \cdot \pi_{\mathrm{FP}}$ (S with no Fast Pass).

Following (Bordalo et al., 2013; Kőszegi and Szeidl, 2013) attention to an attribute increases in its contrast within the choice set, but based on norms retrieved from past experience (Bordalo et al., 2020). Here price contrast has two components, which we formally separate in footnote 10. One is the relative location of a good's price in the menu: a price that stands out from its neighbors draws attention to that good, which is the source of decoy and compromise effects in Bordalo et al. (2013). The other is the overall price variation in the menu: striking or surprising price differences lead the rider to represent the problem as one of "minimizing cost" and to attend to prices across all goods, whereas small differences lead her to attend to other attributes. The second component is the one at work here. Fast Pass is priced far above the incumbents, so its arrival creates a large and surprising price variation on a screen where prices used to be close, and our findings are those of heightened attention to prices across the board rather than of a reallocation among neighboring goods.

To capture this idea, we model price salience as increasing in the perceived price variation, measured as the difference between the *normal* high and low prices in the menu. Let $\overline{p}_r$ be the highest price in the menu $\mathcal{M}_r$ in round r. The associated high-price norm is $\lambda \cdot \overline{p}_r + (1-\lambda) \cdot \sum_{u=0}^{r-1} \overline{p}_u / r$, a linear combination of the highest current price $\overline{p}_r$, which cues recall of similar high past prices, and the average highest price in the database.

In the first Fast Pass session, $r = 1$, the database consists of the pre-experimental period, when Standard was the most expensive option, so $\overline{p}_0 = p_{\mathrm{S}}$. Parameter $\lambda \in (0, 1)$ captures the strength of cued recall (assumed constant for simplicity). The low-price norm is always 0, the price of the outside option. As a result, the perceived price variation in round r is the high-price norm itself.[10]

Let $\Delta = p_{\mathrm{F}} - p_{\mathrm{S}} > 0$ be the price difference between Fast Pass and Standard. Given the i.i.d. presence of Fast Pass, the average high-price norm of treated riders in round r, conditional on the current menu, is equal to

$$p_r^n = \begin{cases} p_{\mathrm{S}} + \lambda \cdot \mathbf{1}_r \cdot \Delta, & r = 1, \\ p_{\mathrm{S}} + \left[\lambda \cdot \mathbf{1}_r + (1 - \lambda) \cdot \frac{r-1}{r} \cdot \phi\right] \cdot \Delta, & r > 1, \end{cases} \tag{5}$$

where $\mathbf{1}_r$ is an indicator for Fast Pass availability at r. In the first experimental round, $r = 1$, the norm is higher than the historical high price p_{S} if and only if Fast Pass is available ($\mathbf{1}_1 = 1$). In subsequent rounds, it is higher than p_{S} also because of past experiences with Fast Pass. The database term $(1 - \lambda) \cdot \frac{r-1}{r} \cdot \phi \cdot \Delta$ contains the i.i.d. probability ϕ because we consider the average population norm, but our results extend to short-run heterogeneity in riders' exposure to p_{F}.

Riders' attention to price in round r is equal to $\sigma_r = \sigma(p_r^n)$ for treated riders and is constant at $\sigma_{\mathrm{C}} = \sigma(p_{\mathrm{S}})$ for Control riders, where $\sigma(\cdot)$ is the strictly increasing salience function. Clearly, $\sigma_r \geq \sigma_{\mathrm{C}}$, so price is weakly more salient for treated than control riders. To see the effect on demand, let $d_j(\mathcal{M}, \sigma)$ be the probability that option j is chosen

[10] In Bordalo et al. (2020) the salience of good j's price is an increasing function of its surprise relative to a memory norm, e.g. $(p_j - p_j^n)^2$. In a menu of M goods, a natural extension lets the salience of j's price increase in its average surprise relative to the norms of all goods, $\frac{1}{M}\sum_m (p_j - p_m^n)^2 = (p_j - \mathrm{E}p^n)^2 + \mathrm{Var}(p^n)$, where $\mathrm{E}p^n$ and $\mathrm{Var}(p^n)$ are the mean and variance of the norms in the menu. The first term is the relative location of j, the source of decoy and compromise effects in Bordalo et al. (2013); the second is the price variation in the menu, common to all goods, on which we focus and which we approximate by the distance between the normal highest and lowest prices.

from menu $\mathcal{M}$ under salience σ. For treated riders who are shown Fast Pass in round r ($\mathbf{1}_r = 1$), the effect of adding Fast Pass on the demand for $j = 0, \mathrm{W}, \mathrm{S}$ decomposes as

$$d_j(\mathcal{F}, \sigma_r) - d_j(\mathcal{C}, \sigma_{\mathrm{C}}) = \underbrace{\left[d_j(\mathcal{F}, \sigma_r) - d_j(\mathcal{C}, \sigma_r)\right]}_{\text{standard}} + \underbrace{\left[d_j(\mathcal{C}, \sigma_r) - d_j(\mathcal{C}, \sigma_{\mathrm{C}})\right]}_{\text{non-standard}}. \tag{6}$$

The standard term is the change in demand due to adding Fast Pass at constant attention. It is weakly negative for all incumbent goods. In our specific setting, this effect is negative for S (riders with a high value of time θ switch to F) and zero for 0 and W. The non-standard term instead captures the effect of Fast Pass through the changing price weight. With stable tastes the price weight is constant, so this term is zero. Under salience, adding Fast Pass boosts the price weight, increasing riders' preference for cheaper options. Thus, the term is positive for 0, negative for S, and ambiguous for W.

Positive effects of Fast Pass on the demand for 0 and W means that the non-standard terms dominate the standard ones, rejecting regularity and hence rationality. This mechanism can reproduce our experimental results across four types of treated riders.[11]

Prediction 1: higher first session demand for 0 and W. At $r = 1$, if p_{F} is shown, $p_1^n = p_{\mathrm{S}} + \lambda \cdot \Delta > p_{\mathrm{S}}$. Fast Pass increases price salience $\sigma_1 > \sigma_{\mathrm{C}}$. Thus, demand increases for the outside option, $d_0(\mathcal{F}, \sigma_1) > d_0(\mathcal{C}, \sigma_{\mathrm{C}})$, and for W&S, $d_{\mathrm{W}}(\mathcal{F}, \sigma_1) > d_{\mathrm{W}}(\mathcal{C}, \sigma_{\mathrm{C}})$, if and only if riders switching from Standard to W&S outnumber those who switch from W&S to 0. For a small salience shock, this condition reduces to

$$\pi_S \cdot f(\sigma_1 \cdot \pi_S) > \pi_W \cdot f(\sigma_1 \cdot \pi_W), \tag{7}$$

namely the density of $\log \theta$ is higher at the Standard–W&S margin than at the W&S–non-purchase margin. The condition holds if (7) holds for $\sigma \in [\sigma_{\mathrm{C}}, \sigma_1]$.

[11] With logit demand the standard term is negative also for W and 0, which means that salience effects must be even stronger to produce violations of regularity.

Prediction 2: stronger effect with higher p_{F}. A higher Fast Pass price p_{F} increases the norm and hence price salience σ_r, increasing the magnitude of the above effect on the demand for 0 and, if condition (7) holds, also of the effect on the demand for W.

Salience thus predicts the patterns of Table 1. Memory produces two other patterns.

Prediction 3: stronger effect with repeated exposure. When Fast Pass is shown in round $r > 1$, the norm of treated riders is higher than when it is shown in the first round, $r = 1$, because the database adds to the perceived price contrast, $p_r^n > p_1^n$ for $r > 1$. Thus, subsequent exposure to Fast Pass causes a stronger increase in the demand for 0 and, if condition (7) holds at these higher salience levels, in the demand for W.

Prediction 4: effect at $r > 1$ when Fast Pass is not shown. Even if Fast Pass is not shown, at $r > 1$ the perceived price contrast of treated riders is higher than that of control riders by $(1 - \lambda) \cdot \frac{r-1}{r} \cdot \phi \cdot \Delta$, due to past exposure to Fast Pass. Thus, their demand for 0 is higher relative to control, especially if p_{F} is higher, and the same occurs for W if condition (7) holds at the relevant salience levels. Both effects are, however, smaller than when Fast Pass is shown.

A potential confound is inertia: riders may repeat their previous choice out of habit. Inertia cannot explain the first-session effects or their price gradient (Predictions 1 and 2). It could mimic Predictions 3 and 4, because choices induced by Fast Pass accumulate over sessions and carry over into sessions without Fast Pass. Inertia, however, can only perpetuate a first-session effect, not create one. At the low Fast Pass price, the first-session effect on W&S is zero (-0.02 p.p., Table 1), yet it reaches 0.6 p.p. over all sessions (Table 2). More generally, inertia implies that the all-sessions effect is the same multiple of the first-session effect at every price and for every benchmark. Relative to the low price, the multiple is 2.2, and inertia of this strength implies the 70 percent drop without Fast Pass in Table 3. Relative to control, it is four to five, because of the growth

at the low price. Inertia thus fits the differences across prices but leaves unexplained an effect of about 0.65 p.p. common to all three prices.[12] Memory can generate this effect, because the database grows with every exposure, whether or not the first exposure influenced choice.

4 Conclusion

In a natural field experiment with 7.2 million riders we show that adding a premium product violates regularity, the weakest implication of stable tastes. Contrary to compromise and decoy effects, the premium option pushes riders toward the cheapest option and out of the market. The effect grows with the premium price and with exposure, and survives in sessions where the premium option is absent. These patterns defy rational inference and inattention, and are produced by a model where a striking price, cued on the screen or retrieved from memory, raises attention to prices bottom up.

Three implications follow. First, a firm widening the range of prices on display can make consumers price-conscious rather quality-focused, so that the direct demand for the new product is partly offset by higher price sensitivity. The effect is not small. At the high Fast Pass price, Table 2 shows that for every 100 treated sessions Fast Pass gains 1.1 riders, while 0.3 stop buying. Since an upgrade earns only the markup while a lost ride costs the whole fare, the lost rides alone wipe out a Fast Pass markup around a quarter of the fare. The 1.3 riders downgrading from Standard to W&S add to the loss.

[12] Suppose a rider repeats her previous choice with probability ρ and otherwise chooses like in control with probability $1-\phi$, and increases her W&S probability by δ_p with probability $\phi = 0.33$. In steady state the first-session effect of price arm p relative to any arm q is $(1-\rho)\phi(\delta_p - \delta_q)$ and the all-sessions effect is $\phi(\delta_p - \delta_q)$, so the growth factor $1/(1-\rho)$ is the same for every comparison. Without Fast Pass the effect is $\rho\phi(\delta_p - \delta_q)$, with it $\rho\phi(\delta_p - \delta_q) + (1-\rho)(\delta_p - \delta_q)$, so the drop without Fast Pass is $1/[1+\phi\rho/(1-\rho)]$. The growth of 2.2 relative to the low price (0.32 to 0.7 p.p. at the high price) implies $\rho \approx 0.55$ and a drop of 72 percent, as observed. The same ρ implies all-sessions effects relative to control of 0.4 and 0.7 p.p. at the medium and high price, against 1.1 and 1.3 observed. The conclusion is unchanged with a finite number of sessions per rider.

Second, demand estimation imposing regularity can be problematic: in our experiment, a "regular" demand model would impose that Fast Pass boosts demand for Lyft, attributing the increased shares of W&S and not buying to a negative taste shock. It would thus measure the benefit of Fast Pass from upgrades but miss entirely its price-sensitivity cost. Third, and finally, our findings show the importance of formalizing the cognitive drivers of context effects: they can allow to unveil the conditions in which compromise or anti-compromise effects can arise and the amplification of effects through exposure and cues. Accounting for cognition can help recover hedonics and optimal pricing from observed demand, as in Gennaioli et al. (2026).

References

Andreoni, J. and Miller, J. (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753.

Bertrand, M., Karlan, D., Mullainathan, S., Shafir, E., and Zinman, J. (2010). What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment*. *Quarterly Journal of Economics*, 125(1):263–305.

Bodner, R. and Prelec, D. (1994). The centroid model of context dependent choice. *Unpublished Manuscript, MIT*.

Bordalo, P., Gennaioli, N., Lanzani, G., and Shleifer, A. (2026). A cognitive theory of reasoning and choice. *The Quarterly Journal of Economics*, page qjag023.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and Consumer Choice. *Journal of Political Economy*, 121(5):803–843.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, Attention, and Choice. *The Quarterly Journal of Economics*, 135(3):1399–1442.

Bushong, B., Rabin, M., and Schwartzstein, J. (2021). A model of relative thinking. *The Review of Economic Studies*, 88(1):162–191.

Cai, J. (2026). Contract design and insurance demand. Technical report, National Bureau of Economic Research.

Caplin, A., Dean, M., and Martin, D. (2011). Search and satisficing. *American Economic Review*, 101(7):2899–2922.

Crosetto, P. and Gaudeul, A. (2016). A monetary measure of the strength and robustness of the attraction effect. *Economics Letters*, 149:38–43.

Enke, B., Graeber, T., and Oprea, R. (2023). Confidence, self-selection, and bias in the aggregate. *American Economic Review*, 113(7):1933–1966.

Gennaioli, N., Rodemeier, M., and Sun, G. (2026). Nonlinear pricing and revealed salience. Technical Report 12769, CESifo, Munich.

Gu, X., Kannan, P., and Ma, L. (2018). Selling the premium in freemium. *Journal of Marketing*, 82(6):10–27.

Herne, K. (1999). The effects of decoy gambles on individual choice. *Experimental Economics*, 2(1):31–40.

Huber, J., Payne, J. W., and Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of consumer research*, 9(1):90–98.

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *American Economic Review*, 98(5):2127–2149.

Kivetz, R., Netzer, O., and Srinivasan, V. (2004). Alternative models for capturing the compromise effect. *Journal of marketing research*, 41(3):237–257.

Kőszegi, B. and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly journal of economics*, 128(1):53–104.

Kroll, E. B. and Vogt, B. (2012). The relevance of irrelevant alternatives. *Economics Letters*, 115(3):435–437.

Kuksov, D. and Villas-Boas, J. M. (2010). When more alternatives lead to less choice. *Marketing Science*, 29(3):507–524.

List, J. A. (2003). Does market experience eliminate market anomalies? *The quarterly journal of economics*, 118(1):41–71.

Long, X., Sun, J., Dai, H., Zhang, D., Zhang, J., Chen, Y., Hu, H., and Zhao, B. (2025). The choice overload effect in online recommender systems. *Manufacturing & Service Operations Management*, 27(1):249–268.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98.

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of political economy*, 94(4):796–821.

Miller, K. M., Sahni, N. S., and Strulov-Shlain, A. (2026). Sophisticated consumers

with inertia: Long-term implications from a large-scale field experiment. *American Economic Review*.

Natenzon, P. (2019). Random choice and learning. *Journal of Political Economy*, 127(1):419–457.

Rafai, I., Babutsidze, Z., Delahaye, T., Hanaki, N., and Acuna-Agost, R. (2022). No evidence of attraction effect among recommended options: A large-scale field experiment on an online flight aggregator. *Decision Support Systems*, 153:113672.

Rodemeier, M. (2025). Buy baits and consumer sophistication: Field evidence from instant rebates. *American Economic Journal: Economic Policy*, 17(2):30–59.

Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. In *The Foundations of Price Theory Vol 4*, pages 101–116. Routledge.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of consumer research*, pages 158–174.

Simonson, I. and Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of marketing research*, 29(3):281–295.

Somerville, J. (2022). Range-dependent attribute weighting in consumer choice: An experimental test. *Econometrica*, 90(2):799–830.

Sonsino, D. (2010). The irrelevant-menu affect on valuation. *Experimental Economics*, 13(3):309–333.

Tirole, J. (1988). *The theory of industrial organization*, volume 1. MIT press Cambridge, MA.

Tversky, A. and Simonson, I. (1993). Context-dependent preferences. *Management science*, 39(10):1179–1189.

Wernerfelt, B. (1995). A rational reconstruction of the compromise effect: Using market data to infer utilities. *Journal of Consumer Research*, 21(4):627–633.

Woodford, M. (2012). Inattentive valuation and reference-dependent choice. Working paper, Columbia University.