POSTED: Feb 28, 2014

On the Equilibrium Effects of Nudging*

Ran Spiegler†

January 15, 2014

Abstract

Consumers' decision biases make them vulnerable to market exploitation. "Libertarian paternalism" (a.k.a "nudging") is the viewpoint that this problem can be mitigated by "soft" interventions like disclosure or "default architecture". However, the case for nudging is often made without an explicit model of the boundedly rational choice procedures that lie behind consumer biases. I demonstrate that once such models are incorporated into the analysis, equilibrium market reaction to nudges can reverse their theoretical consequences.

1 Introduction

Our everyday thinking about consumer protection relies on some intuitive, informal notion of bounded rationality. In extreme cases like duress or dementia, the presumption is that the consumer cannot make an intelligent

*This paper has enjoyed financial support from the Pinhas Sapir Center for Development. It grew out of a project that was supported by ERC Grant no. 230251. I thank Kfir Eliaz for useful discussions.

†Tel Aviv University, University College London and CFM. URL: http://www.tau.ac.il/~rani. E-mail: r.spiegler@ucl.ac.uk.

choice and would accept manifestly exploitative contracts. In these cases, we have a strong intuition that some kind of protection is in order, including extreme measures such as voiding the contract.

The more interesting cases for an economist are those where the consumer is a "reasonable" decision maker who is capable of making free and intelligent choices. However, a "reasonable" decision maker is not an infallible or infinitely rational one; he is naturally limited in his computational abilities and the quality of his understanding of the market environment; he often suffers from attention deficits due to the many environments he interacts with; and being human, he occasionally succumbs to temptation, self-delusion or wishful thinking. Since this makes him susceptible to market exploitation, the challenge for economists and legal scholars is how to think about consumer protection in this context. In particular, can regulators minimize market exploitation without resorting to measures that impose limits on contractual freedom?

One school of thought is encapsulated in the words of a former FTC chairman: "robust competition is the best single means for protecting consumer interests" (Muris (2002)). However, a key message of the theoretical literature on "behavioral industrial organization" (see Ellison (2006), Armstrong (2008) and Huck and Zhou (2011) for surveys, and Spiegler (2011) for a graduate-level textbook) is that on the whole, models of market competition among profit-maximizing firms over boundedly rational consumers do *not* support this motto. For instance, Spiegler (2006) examined a market model in which firms compete in random prices and consumers evaluate them with a simplifying "sampling" procedure, and showed that increasing the number of competitors intensifies obfuscation (in the sense of increasing the variance of the equilibrium pricing strategy), without changing the expected equilibrium price. At some level, this should not be too surprising. After all, if bounded rationality means that consumers are not that good at reaching "correct" decisions, we should not expect that expanding the set

of market options will necessarily make them better off, especially when the complexity of making sound choices is influenced by the firms' endogenous obfuscation tactics.

In the last decade, a new school of thought emerged (Thaler and Sunstein (2003), Camerer et al. (2003), Bar-Gill (2012)), arguing for "soft" interventions that mitigate market exploitation of boundedly rational consumers, while imposing virtually no constraints on contractual freedom. For instance, the regulator can manipulate features of consumers' choice set that would be irrelevant for a rational decision maker, such as the order in which alternatives are presented to the consumer, or the specification of a default alternative. The regulator could also impose user-friendly disclosure requirements that would mitigate misleading or confusing contracts. Choice architecture and disclosure are indeed the prime examples of this approach, dubbed "*libertarian paternalism*" or "*nudging*", which received a major boost in the hugely popular manifesto by Thaler and Sunstein (2008).

The "nudgniks" have been criticized on various philosophical grounds. For instance, it can be argued that any governmental design of the choice set that manipulates consumer behavior goes against libertarian values, especially when it springs from a paternalistic pretense to know consumers' "true preferences". Nevertheless, it is obvious that nudging hurts freedom of choice to a much lesser degree than outright paternalistic measures such as sin taxes or banning contracts. Thaler and Sunstein themselves respond eloquently to many of these criticisms.

This paper offers a different critique, which targets two characteristics of most discussions of "nudging". First, the literature tends to ignore the *equilibrium* market response to libertarian-paternalistic interventions. Second, it typically regards consumers' decision errors and biases as primitive behavioral phenomena, namely "black boxes". In particular, Bar-Gill (2012) and Mullainathan et al. (2012) employ models in which a rational decision involves trading off costs and benefits according to certain decision weights,

such that decision biases are simply modeled as using the wrong weights. Although Mullainathan et al. (2012) allow the weights to be a function of "nudges", they leave this function unspecified. Such "reduced-form" models of consumer biases do not tell an explicit story about where the wrong weights come from, and offer little guidance as to how they could be affected by the obfuscation tactics employed by firms or by regulatory interventions. The two characteristics are thus interrelated: the "reduced-form" approach to modeling consumer biases limits the scope of analyzing the equilibrium effects of "nudging".

Following a different strand in the literature on bounded rationality as a whole (Rubinstein (1998)) and behavioral industrial organization in particular (Spiegler (2011)), I challenge this viewpoint, by considering a sequence of market models in which profit-maximizing firms compete for boundedly rational consumers. In all these models, consumers commit decision errors that intuitively call for consumer protection measures, and there is a natural case for libertarian-paternalistic measures like "default architecture" or disclosure. The difference from the reduced-form approach is that these models involve *explicit boundedly rational choice procedures*, which generate the decision errors of interest. I then analyze the firms' equilibrium response to libertarian-paternalistic interventions, taking these explicit models of consumer choice into account. In each case, I show that the typical libertarian-paternalistic argument can be reversed.

This paper is not the first to make this point. Kamenica et al. (2011) presented an example demonstrating that theoretically, a certain kind of usage-based disclosure may have an adverse effect on consumer welfare, once market equilibrium effects are taken into account. Likewise, Piccione and Spiegler (2012) showed that harmonizing description formats may have a similar counterproductive effect on consumer welfare. Moreover, at a certain level the critique is obvious to economists: the case for equilibrium analysis of "structural models" dates at least as early as Lucas (1976), who made

it in the context of macroeconomic policy. The contribution of this paper is that it articulates the critique in a more systematic and comprehensive manner, presenting new examples and highlighting considerations that are specific to consumer protection (such as the possible tension between the ultimate objective of enhancing consumer welfare and intermediate goals like minimizing obfuscation).

Other recent works have addressed theoretical implications of behavioral industrial organization for market regulation. Heidhues and Koszegi (2010) analyze the welfare implications of banning late fees (an obvious example of "hard" paternalism) in a credit-market model where consumers mispredict their future self-control. Grubb (2012) constructs a dynamic consumption model with limited consumer monitoring of their own past consumption, and discusses various regulatory measures that address the "bill shock" problem that arises in such a model. Armstrong and Vickers (2012) discuss regulation of contingent charges such as overdraft fees in a model with diversely attention consumers. De Clippel et al. (2013) analyze in a model in which consumers optimally allocate their limited attention to individual markets according to the observed prices charged by market leaders. They show that improving consumers' attention weakens market leaders' incentive not to stand out as being too expensive, and this in turn softens competitive pressures and lowers consumer welfare in equilibrium.

2 Default Architecture

Design of default options is arguably the most influential idea in the "nudge" literature. It is based on the observation that decision makers tend to stick to default options in a variety of contexts, from voluntary organ donation to retirement savings. This has given rise to the suggestion, eloquently articulated by Thaler and Sunstein (2008), that policy makers can and should raise participation rates in such settings by appropriate design of default options.

In particular, it has been argued that a switch from "opting in" (a regime in which the default is the outside option) to "opting out" (a regime in which the default is one of the market alternatives) will raise participation rates.

Thaler and Sunstein (2008) acknowledged that default bias is not a primitive phenomenon, and that it originates from more basic considerations, such as loss aversion, limited attention or choice complexity. However, neither they nor others integrated these considerations explicitly into an equilibrium analysis of default design. In this section I attempt such an exercise, in the context of a market in which profit-maximizing firms interact with consumers who follow a natural choice procedure that generates default bias. I show that such analysis may overturn the consequences of default design.

The model extends the formalism of Piccione and Spiegler (2012) by adding an outside option, thus enabling a rich discussion of default architecture. Consider a market that consists of two profit-maximizing firms and a measure one of consumers. The firms offer a homogenous product at zero cost. Each firm $i = 1, 2$ chooses simultaneously a pair (p_i, x_i), where $p_i \in [0, 1]$ is the price of the firm's product, and $x_i \in \{a, b\}$ is the "description format" the firm uses to present the price.

Each consumer chooses according to the following rule. If he can make a price comparison, he selects the cheapest firm. (The tie-breaking rule in case $p_1 = p_2$ is immaterial. However, for completeness, assume the following: the consumer chooses each firm with probability $\frac{1}{2}$ when his default is the outside option, and he sticks to his default when it is one of the firms.) If he cannot make a price comparison, he sticks to his default option. Each firm plays the role of a default option for a fraction $\frac{\alpha}{2}$ of the consumer population, and the outside option (no participation) serves as a default option for the remaining fraction $1 - \alpha$ of the consumer population. A higher value of α represents a move from an "opt in" default design to "opting out".

Comparability of the two market alternatives is entirely a function of the description formats that the firms adopt. Specifically, $\pi(x_1, x_2)$ is the

comparison probability when the profile of formats is (x_1, x_2). In all the examples I will consider, π is a symmetric function: $\pi(x, y) \equiv \pi(y, x)$. The resulting payoff function in the game between the firms is as follows. For every strategy profile $((p_1, x_1), (p_2, x_2))$, firm i's profit is $p_i \cdot s_i((p_1, x_1), (p_2, x_2))$, where s_i is firm i's market share, given by

$$s_i((p_1, x_1), (p_2, x_2)) = \begin{cases} \frac{\alpha}{2} + (1 - \frac{\alpha}{2})\pi(x_1, x_2) & if \quad p_i < p_j \\ \frac{\alpha}{2} & if \quad p_i = p_j \\ \frac{\alpha}{2}(1 - \pi(x_1, x_2)) & if \quad p_i > p_j \end{cases}$$

Note that in this model, consumers do not have an understanding of equilibrium regularities, and so they cannot draw any inferences regarding prices from the observed description formats.

This choice procedure has the key feature that when consumers face a complex choice problem - namely, a situation in which they are unable to make a price comparison - they are reluctant to make an active decision, and instead choose by default (which can also be viewed as procrastinating). The notion that default bias is exacerbated by complexity of the choice problem is supported by various studies (e.g., for "field" evidence on retirement savings, see Iyengar et al. (2004) and Beshears et al. (2010)).

Let us begin with the following simple comparability structure, denoted by π^u: $\pi(x_1, x_2) = 1$ (0) if $x_1 = x_2$ ($x_1 \neq x_2$). The payoff function now has a very simple description: if $x_1 \neq x_2$, firm i earns $p_i \cdot \frac{\alpha}{2}$; and if $x_1 = x_2$, firm i earns $p_i \cdot [\frac{1}{2} + \frac{1}{2} sign(p_j - p_i)]$. The interpretation is that a and b represent different *measurement units*, and consumers can make an active comparison only when prices are denominated in/per the same units. For instance, interest rates on loans can be stated for various time units, and knowing how to compare them requires minimal financial numeracy (in particular, an understanding of compounding).

Comment. I interpret the upper bound on the firms' prices as the consumers' willingness to pay for the product. This creates some tension with the limited

comparability story under the "opt out" rule: if consumers are unable to make a comparison, what prevents firms from raising the price above 1? One answer is that consumers are able to cancel their purchase ex-post, and this ex-post participation constraint prevents over-pricing. Another answer is that consumers understand the price associated with their default option, and since they switch away only when making an active comparison based on a correct price comparison, they will never pay more than 1.

It can be easily shown that the game between the two firms has no pure-strategy Nash equilibrium. Also, when $\alpha = 0$, Nash equilibrium implies that $p_1 = p_2 = 0$, but the firms' format strategies (hence the equilibrium market participation rate) are entirely indeterminate. I will therefore analyze symmetric mixed-strategy equilibria for $\alpha > 0$. A mixed strategy in this model is a joint probability distribution over prices and formats.

Proposition 1 *Let $\alpha > 0$. Under the comparability structure π^u, there is a unique symmetric Nash equilibrium. Firms play a and b with probability $\frac{1}{2}$ each, and independently mix over prices according to the cdf*

$$F(p) = 1 - \frac{\alpha}{2}\left(\frac{1}{p} - 1\right)$$

defined over the interval $[\frac{\alpha}{2+\alpha}, 1]$.

The probability of comparison in equilibrium is $\frac{1}{2}$ for all realized price pairs (p_1, p_2) and every value of α. The latter means that equilibrium choice complexity is invariant to the default architecture. As a result, the equilibrium rate in which consumers switch away from their default option is always $\frac{1}{2}$ ($\frac{1}{4}$) for consumers who are initially assigned to the outside option (one of the firms). The overall rate of market participation is thus $\alpha + \frac{1}{2}(1 - \alpha)$. Since the gross value of the product to consumers is 1, this is also the total surplus in equilibrium.

To calculate firms' equilibrium profits, note that the price $p = 1$ is in the support of the equilibrium pricing strategy. The clientele of a firm that charges this price consists purely of the consumers who were initially assigned to the firm and failed to make a price comparison. The size of this clientele is thus $\frac{\alpha}{4}$, hence each firm earns an equilibrium payoff of $\frac{\alpha}{4}$. As we move from "opt in" to "opt out", equilibrium prices and industry profits go up. The intuition is that under "opt out", firms benefit from consumers' default bias and exploit it through higher prices.

It follows that equilibrium consumer surplus is

$$\left[\alpha + \frac{1}{2}(1-\alpha)\right] - \left[2 \cdot \frac{\alpha}{4}\right] = \frac{1}{2}$$

Thus, equilibrium consumer welfare is *invariant to the default architecture.* As we shift from "opting in" to "opting out", the gain in consumer welfare due to increased market participation is exactly offset by their welfare loss due to the firms' exploitation of default bias. This invariance turns out to be an artefact of the comparability structure.

To demonstrate this point, consider an alternative comparability structure, denoted π^c and defined as follows: $\pi^c(a,a) = 0$, $\pi^c(b,b) = 1$, $\pi^c(a,b) = \frac{1}{2}$.The interpretation is that a is an inherently "complex" format that inhibits comparison (e.g., using technical jargon), whereas b is a "simple" format that facilitates comparison (e.g., using lay terms). Such a structure was analyzed by Chioveanu and Zhou (2013), in a model with $n \geq 2$ firms but without an outside option.

Proposition 2 *Let $\alpha > 0$. Under π^c, the game has a unique symmetric Nash equilibrium. With probability $\frac{\alpha}{2}$, firms play $(p,x) = (1,a)$. With the remaining probability $1 - \frac{\alpha}{2}$, they play $x = b$ and mix over prices according to the cdf*

$$F(p) = \frac{1}{4(2-\alpha)}\left[\alpha^2 - 2\alpha + 8 - \frac{\alpha^2 + 2\alpha}{p}\right]$$

defined over the interval

$$[\frac{\alpha^2+2\alpha}{\alpha^2-2\alpha+8}, 1)$$

Thus, unlike the case of π^u, in symmetric equilibrium under π^c there is a clear correlation between the firms' price and format decisions. When they use the complex format a, they charge the monopoly price $p = 1$. In contrast, when they use the simple format b, they mix over prices below the monopoly level.

The fraction of consumers who make a price comparison in equilibrium is

$$(1-\frac{\alpha}{2})^2 + 2 \cdot \frac{1}{2} \cdot \frac{\alpha}{2}(1-\frac{\alpha}{2}) = 1 - \frac{\alpha}{2}$$

The equilibrium switching rate is thus $1 - \frac{\alpha}{2}$ for consumers who are initially assigned to the outside option, and half this rate for consumers who are initially assigned to one of the firms. We can see that unlike the case of π^u, equilibrium choice complexity is sensitive to the default architecture. The equilibrium market participation rate is

$$\alpha + (1-\alpha)(1-\frac{\alpha}{2}) = \frac{1}{2}\alpha^2 - \frac{1}{2}\alpha + 1$$

The first term on the L.H.S is simply the fraction of consumers who are initially assigned to one of the firms; the second term represents the fraction of consumers who are initially assigned to the outside option but end up making a price comparison. Note that this expression is U-shaped w.r.t α: it attains the maximum of 1 both at $\alpha = 1$ and in the $\alpha \to 0$ limit, and it attains a minimum of $\frac{7}{8}$ at $\alpha = \frac{1}{2}$.

As in the case of π^u, firms' equilibrium format strategies are indeterminate when $\alpha = 0$. Thus, it is instructive to examine the limit of equilibria as $\alpha \to 0$: when firms play $(p, x) = (1, b)$ with probability one. By comparison,

when $\alpha = 1$ (an extreme "opt out" regime), industry profits are maximized and the probability of active choices (and hence the switching rate) attains the minimum.

To calculate consumer welfare in equilibrium, note first that equilibrium industry profits are $\frac{1}{4}\alpha^2 + \frac{1}{2}\alpha$. To see why, note that when a firm plays $(1, a)$, its market share is $\frac{\alpha}{2} \cdot (\frac{\alpha}{2} + \frac{1}{2}(1 - \frac{\alpha}{2})) = \frac{1}{8}\alpha^2 + \frac{1}{4}\alpha$. Since the product's value for consumers is 1, consumer surplus is obtained by subtracting industry profits from the participation rate:

$$\begin{aligned} & [\frac{1}{2}\alpha^2 - \frac{1}{2}\alpha + 1] - [\frac{1}{4}\alpha^2 + \frac{1}{2}\alpha] \\ = \; & 1 + \frac{1}{4}\alpha^2 - \alpha \end{aligned}$$

This expression is clearly decreasing in α. In particular, the limit equilibrium for $\alpha \to 0$ is unambiguously the optimal one as far as consumers are concerned. There is full market participation and the price is competitive. Thus, the optimal default design is a slightly perturbed "opt in" (the perturbation is needed to eliminate other equilibria).

The intuition is that in this model, default bias results from the complexity of making price comparisons. Under π^c, the perturbed "opt in" policy encourages firms to maximize comparability (by using the simple format), because firms only get consumers thanks to active choices. Ease of comparison in turn strengthens the price competition, and this leads to a competitive outcome. In contrast, under an "opt out" policy, all consumers participate, but firms have a weaker incentive to maximize comparability because consumers stick to their default option when they are unable to make a comparison. This in turn raises equilibrium prices and reduces consumer switching. In contrast, under π^u there is no distinction between simple and complex formats: when one firm plays each format with probability $\frac{1}{2}$, the comparison probability is $\frac{1}{2}$, regardless of the format employed by the other firm.

The analysis in this section has a straightforward interpretation in terms

of regulating *auto-renewals* in settings like insurance or magazine subscription. Think of "opting out" as auto-renewal and of "opting in" as a regime in which auto-renewals are not allowed. Then, intuitively, when auto-renewals are banned, firms have a stronger incentive to make prices transparent, which in turn impels them to offer more competitive prices. The model confirms the intuition that in this environment, banning auto-renewals may be superior in terms of consumer welfare. However, the confirmation is qualified, because as we saw, this depends on the comparability structure.

Let us return to the retirement savings setting, which features prominently in the literature on default architecture. In the case of 401(k) accounts, we do not have a market in which firms directly and freely compete for consumers, but a market that is effectively regulated by the employers, who mediate the transaction between employees and investment funds. The exercise in this section suggests that if this de-facto regulation were lifted, pure "opt out" default architecture could theoretically harm consumers due to equilibrium effects.

This section has demonstrated the importance of accounting for the procedural origins of consumers' decision biases. Default bias is not a primitive phenomenon; it is often a consequence of the boundedly rational procedures that consumers employ in response to complex choice problems. I focused on the effect of limited comparability between price formats on consumer inertia, and showed that a change in consumers' default option may affect the firms' incentive to facilitate or obstruct price comparisons, and this in turn affects competitive pressures and therefore consumer welfare.

Comment: Harmonizing formats

One advantage of thinking about consumer biases in terms of underlying choice procedures is that it enables us to thinking about a variety of consumer protection measures. Consider the model of this section with an "opt out" default rule ($\alpha = 1$). Modify the comparability structure π^c such that $\pi^c(b,b) < 1$ (while continuing to assume $\pi^c(a,a) = 0$, $\pi^c(a,b) = \frac{1}{2}$). This

modification clearly weakens comparability. It can be interpreted as the *opposite* of a measure that is intuitively beneficial to consumers, namely harmonizing different description formats.

Nevertheless, Piccione and Spiegler (2012) show that such a change in π^c would induce *lower* expected equilibrium profits, and consequently higher consumer welfare. The implication is that the opposite intervention of harmonizing formats can have *adverse* effects for consumers. The intuition behind this result is as follows. Equilibrium industry profits are determined by the payoff that the strategy $(p, x) = (1, a)$ generates:

$$\frac{\alpha}{2}\left[\lambda + (1-\lambda)(1-\pi^c(a,b))\right] = \frac{1}{2}\left[\lambda + \frac{1}{2}(1-\lambda)\right]$$

where λ is the probability that the equilibrium strategy assigns to $(1, a)$. The change in π^c does not have any direct effect on this expression. However, it raises the market share of a firm playing $(1-\varepsilon, b)$, as long as $\varepsilon > 0$ is sufficiently small. To restore equilibrium, λ must go down, and this means that the equilibrium outcome is more competitive. This is another demonstration that a consumer-protection measure that intuitively benefits boundedly rational consumers can harm them, once the equilibrium market response is incorporated into the theoretical analysis.

3 Product-Use Disclosure

Another "libertarian paternalistic" intervention that has received considerable attention is regulating disclosure (Thaler and Sunstein (2008) and Bar-Gill (2012) are usefully comprehensive references). The literature distinguishes between disclosure of product use and disclosure of product attributes. The former aims to facilitate the evaluation of a complex price plan according to an estimated level of consumption (on the basis of the consumer's own past behavior, or the behavior of other consumers in simi-

lar circumstances). The latter aims to correct biases and omissions in the perceived salience of various attributes.

This section focuses on product-use disclosure in the context of a market model based on DellaVigna and Malmendier (2004). The presentation is borrowed from Spiegler (2011, Ch. 3). Two firms provide an identical service and compete for a single consumer. Each firm i simultaneously commits to a non-linear price scheme $t_i : [0, \infty) \to \mathbb{R}$, where $t_i(x)$ is the payment the consumer makes to firm i conditional on selecting this firm and subsequently choosing the consumption level x. Following DellaVigna and Malmendier (2004), I restrict the price scheme t_i to take the form of a two-part tariff, namely $t_i(x) = A_i + p_i x$. Both firms face the same marginal cost $c = \frac{1}{4}$. The firm's profit is zero if it is not chosen by the consumer, and $t(x) - cx$ if the consumer chooses the firm and proceeds to consume x.

Consumers have dynamically inconsistent preferences. From their ex-ante point of view, their willingness to pay for any quantity x is u, given by $u(x) = \ln(1 + \min\{x, \varepsilon\})$, where $\varepsilon \in (0, 1)$. The interpretation is that there is a relatively small quantity ε that the consumer needs, and he deems any additional consumption as superfluous from an ex-ante perspective. After the consumer accepts a contract but before he selects his consumption level, he starts attaching value to additional consumption, and his willingness to pay changes into $v(x) = \ln(1 + x)$ for all $x \geq 0$. Various forces can generate this pattern of changing tastes. For instance, u may reflect a "cold state" evaluation, while v represents a "hot state" taste for immediate gratification. Stories of this kind are relevant for credit cards, smart-phone applications, etc.

It is conventional in the literature to distinguish between "sophisticated" consumers who correctly anticipate the future change in their tastes and "naive" consumers who erroneously believe that their preferences will not change. I assume that the consumer is *naive* - that is, he believes that if he selects firm i, he will proceed to choose x to maximize $u(x) - t_i(x)$, whereas

in fact he will choose x to maximize $v(x) - t_i(x)$. Let us examine symmetric Nash equilibrium in the game between the firms (the focus on symmetry is merely to simplify exposition).

Proposition 3 ((DellaVigna and Malmendier (2004))) *In symmetric Nash equilibrium, each firm offers a two-part-tariff t^* characterized by $A^* = -\frac{1}{4}$ and $p^* = \frac{1}{2}$. This scheme induces a consumption quantity $y^* = 1$.*

The equilibrium per-unit price is above the marginal cost. Competitive pressures push the firms' actual equilibrium profits to zero, which means that the lump-sum payment A^* is negative in equilibrium. That is, when accepting t^*, the consumer expects to consume the small quantity ε (because $p^* < u'(x)$ for $x < \varepsilon$), and his main attraction is the relatively large negative lump-sum payment. However, the consumer's equilibrium from the ex-ante perspective (but given his actual ex-post behavior) is $u(y^*) - A^* - p^*y^* = \ln(1+\varepsilon) + \frac{1}{4} - \frac{1}{2} \cdot 1 = \ln(1+\varepsilon) - \frac{1}{4} < 0$. That is, the consumer ends up being exploited (from an ex-ante point of view) despite market competition.

Imagine that in response to this situation, a regulator introduces usage-based disclosure. Specifically, for each contract offered in the market, the regulator mandates the disclosure of the *effective average price given the historical consumption quantity.* Thus, if the historical consumption quantity is some x^*, each offered two-part tariff t will be accompanied by the disclosure of the *effective average price* $t(x^*)/x^*$. Assume that the consumer "obediently" chooses the firm that offers the contract with the lowest disclosed effective average price (with a symmetric tie-breaking rule). Having selected a contract t, the consumer proceeds to choose the consumption quantity that maximizes $v(x) - t(x)$. The interpretation is that in each period, there is a new generation of consumers who make a once-and-for-all decision, and the disclosure is informed by the historical behavior of previous consumer generations.

It should be emphasized that introducing product-use disclosure into a DellaVigna-Malmendier model is not an arbitrary move; one of the key applications of the model was indeed to credit markets, a context in which usage-based disclosure is commonly discussed. In particular, Bar-Gill (2012) presents a DellaVigna-Malmendier model in his discussion of credit cards, and recommends product-use disclosure as a potential remedy for over-consumption patterns of the kind captured by Proposition 3. However, Bar-Gill does not corroborates this recommendation with an equilibrium analysis of such an intervention in the context of the DellaVigna-Malmendier model.

This extension requires us to modify the notion of a stable market outcome. We will say that the pair (t^*, x^*) is *stable* if the following conditions hold: (i) $x^* = \arg\max_x[v(x) - t^*(x)]$; (ii) no firm has an incentive to deviate from t^* to another contract t, given the consumer's rule for choosing a contract and his choice of consumption quantity under his selected contract. Condition (i) means that in order for the consumption quantity x^* to persist, it must be optimal for consumers (according to their *ex-post* preferences) given the equilibrium contract t^*. Condition (ii) reflects the notion that if x^* is a stable consumption level, it becomes the historical quantity that informs the calculation of the effective average price. If a firm deviates from t^* to some other contract t, then either the consumer will not choose t because it does not have a lower disclosed effective average price, or t does attract the consumer (because $t(x^*)/x^* < t^*(x^*)/x^*$) and yet the firm does not make a higher profit given the way the consumer actually chooses under t (i.e., $t(x^{**}) - cx^{**} \leq t(x^*) - cx^*$, where $x^{**} = \arg\max_x[v(x) - t(x)]$.

Proposition 4 *There is a unique stable pair* (t^*, x^*)*, where* $t^*(x) = \frac{1}{4}x$ *for all* x*, and* $x^* = 3$*.*

Thus, on the face of it, product-use disclosure is effective in the sense that it induces a competitive outcome. Equilibrium pricing is reduced to linear pricing, with marginal-cost per-unit pricing and no lump-sum payments.

However, this competitiveness is defined in terms of v, the consumer's ex-post preferences, rather than in terms of his ex-ante preferences. In other words, the firms are encouraged to compete for "the wrong self". Note that $x^* > y^*$ - that is, the consumer's consumption level is higher than before the intervention, because the per-unit price has decreased. Moreover, the consumer's ex-ante utility is $u(x^*) - A^* - p^*x^* = \ln(1+\varepsilon) - \frac{3}{4}$. That is, the intervention has *lowered* the consumer's ex-ante utility.

In fact, in terms of consumer behavior and firms' profits, the equilibrium outcome is the same as the one that would emerge in Nash equilibrium of the original game - prior to the intervention - if firms could use *any* non-linear pricing scheme, rather than being restricted to two-part tariffs (see Spiegler (2011, Ch. 2)). Thus, the intervention ends up simulating an environment in which firms can use any non-linear price scheme, and this only intensifies market exploitation. Once again, we see that an intervention that purports to address a market failure due to a consumer bias ends up harming consumers, once equilibrium effects and the choice procedure underlying the bias are taken into account.

4 Product-Attribute Disclosure

Many products and services have multiple price and quality attributes. Consumers often neglect some of these attributes because they are less salient than others. For instance, when thinking about the actual cost of a loan, borrowers may pay more attention to the basic interest rate than to the late fees. Consumer contracts often have a headline price, which is salient, as well as qualifications that appear in small print. Some products have future add-ons that the consumer may fail to take into account at the time he chooses the product (e.g. replacement ink cartridges in printers). The question is whether mandated disclosure of the less salient product attributes can help consumers making better decisions in this context.

Note that the implicit assumption behind attribute disclosure is that once the consumer is fully aware of all attributes, he will execute a rational evaluation. However, suppose that the consumer has a fixed "*attention budget*", in the sense that he can only take into consideration a subset of attributes; is it now obvious that attribute disclosure will make him better off? Rather than increasing the consumer's attention budget, disclosure might simply *reallocate* it among the various attributes.

The tendency to focus on a small subset of attributes can also arise from an aversion to performing difficult trade-offs. For instance, when evaluating retirement saving plans, how does one trade off the bequest motive with ensuring a decent living standard in old age? This is a difficult trade off, both cognitively and emotionally, and a natural response is to "forget" some of the relevant attributes in order to simplify the act of choice. In this case, too, attribute disclosure may change the likelihood that any individual attribute is neglected, but not necessarily the overall tendency to neglect attributes.

In this section I use a new model due to Bachi and Spiegler (2014) to capture these considerations. Unlike the previous sections, here none of the results will be new; the only novel contribution is my interpretation of the results in terms of attribute disclosure. Therefore, I report the results briefly and refer the reader to the original paper for more general statements of the results and their proofs.

Our market will (once again) consist of two firms and a measure one of consumers. Each firm $i = 1, 2$ simultaneously chooses a product that is fully characterized by a quality vector $(q_i^1, q_i^2) \geq (0, 0)$. Firm i's profit conditional on being chosen is $1 - \frac{1}{2}(q_i^1 + q_i^2)$. This quantity, multiplied by the measure of consumers who choose the firm, constitutes its overall profit. I refer to $\bar{q}_i = \frac{1}{2}(q_i^1 + q_i^2)$ as the "true quality" of the firm's product. Conventionally rational consumers would be endowed with some strictly increasing and continuous function $u(q^1, q^2)$, and they would always choose the firm that sells

the highest-u product. In Nash equilibrium, firms would offer quality vectors that maximize u subject to the zero-profit condition (i.e., $\bar{q} = 1$).

Now suppose that quality dimension 2 is "shrouded", such that all consumers focus entirely on dimension 1 and choose the firm that offers the highest quality along this dimension (with symmetric tie-breaking). In this case, the model turns into a special case of Gabaix and Laibson (2006)).

Proposition 5 (Gabaix and Laibson (2006)) *When consumers choose entirely according to dimension* 1*, the game between the two firms has a unique Nash equilibrium: each firm plays* $(q^1, q^2) = (2, 0)$.

Thus, when consumers focus entirely on dimension 1, competition is effectively restricted to this dimension, and this enables firms to choose the lowest possible quality along dimension 2. Competitive pressures drive quality along dimension 1 up until firms get zero profits. In terms of average quality, equilibrium products are the same as in the case of conventionally rational consumers. However, they are "misleading", in the sense that the quality that each consumer perceives according to the dimension he focuses on ($q^1 = 2$) is higher than the true average quality $\bar{q} = 1$. In fact, the equilibrium strategy *maximizes* $q^1 - \bar{q}$ subject to the constraint that firms earn non-negative profits - in this respect, it is "maximally obfuscating".

Imagine that a regulator responds to this state of affairs by mandating disclosure that "unshrouds" dimension 2, and that the intervention is successful in the sense that it makes both dimensions equally salient. However, as suggested earlier, suppose that this does not turn consumers into rational "trade-off machines". Instead, it merely reallocates the consumers' "attention budget" between the two dimensions, such that every consumer focuses on a *uniformly drawn* single dimension. As a result, when $q_i^k > q_j^k$ for both $k = 1, 2$ (that is, firm i's product strictly dominates firm j's product, the consumer chooses firm i; but when $q_i^1 > q_j^1$ and $q_i^2 < q_j^2$, each firm gets

half the consumer population. What are the equilibrium implications of this intervention?

Proposition 6 (Bachi and Spiegler (2014)) *When each consumer chooses according to a uniformly drawn single dimension, the game between the two firms has a unique symmetric Nash equilibrium: each firm chooses* $q^1+q^2=1$ *with probability one, and draws* q^1 *uniformly from* $[0,1]$.

Thus, in equilibrium, firms offer products of true average quality $\bar{q}=\frac{1}{2}$, but the breakdown into the two dimensions is random. Note that no market alternative ever dominates the other market alternative in equilibrium - that is, consumers always face hard choices in equilibrium. As to the equilibrium amount of obfuscation, the difference between perceived and true quality (which is $q^1-\bar{q}$ and $q^2-\bar{q}$ with probability $\frac{1}{2}$ each) is uniformly distributed over $[-\frac{1}{2},\frac{1}{2}]$, compared with the deterministic gap $q^1-\bar{q}=1$ prior to the intervention. Thus, the amount of obfuscation is lower, both in real and absolute terms.

This exercise provides another demonstration that having an explicit "story" behind observed consumer biases matters for the equilibrium analysis of nudging. When consumers seem to be ignoring certain attributes, we should ask whether this is a manifestation of simple unawareness, or a result of deeper psychological forces, such as intrinsic attention deficit or aversion to trade-offs. When the latter is the case, consumers will continue to ignore product attributes even if the regulator mandates disclosure, but the neglected attributes will be less predictable, giving rise to a less competitive market outcome.

Comment: proximate vs. ultimate objectives

The ultimate objective of consumer protection policies is to improve consumer welfare. Since consumer welfare is hard to measure, regulators often

pursue easier-to-measure proximate objectives that are presumed to be positively correlated with consumer welfare. In particular, when the rationale behind consumer protection is to address market exploitation due to consumers' decision biases, there is a strong intuition that the amount of obfuscation that we observe is negatively correlated with consumer welfare, and therefore a proximate objective should be to reduce obfuscation.

Let us reexamine this intuition in light of the present model. We saw that attribute disclosure, interpreted as a shift from a maximally asymmetric allocation of the consumer's "attention budget" to a maximally symmetric one, leads to a combination of lower true quality, coupled with a lower amount of obfuscation (defined in terms of the gap in real or absolute terms between perceived and true quality). Thus, there is tension between the ultimate objective - maximizing consumer welfare (defined in terms of the average quality the consumer gets in equilibrium) - and the proximate objective, namely reducing the amount of obfuscation.

This tension is not shared by all "behavioral industrial organization" models. In Spiegler (2006), regulatory interventions like increasing the number of competitors can lead to lower consumer surplus in equilibrium, coupled with intensified obfuscation (measured by the expected gap in absolute terms between perceived and true net utility - the expected gap in real terms is always zero in that model). The lesson is that analyzing consumers' decision errors in terms of their procedural origins enables us to examine whether intuitive proximate criteria for consumer protection (e.g. minimizing obfuscation) are in principle consistent with the ultimate objective of maximizing a social welfare function.

Default architecture revisited

The model of this section offers a different approach to limited comparability in markets, compared with the model of Section 2. Here, limited comparability is not determined by description formats that are independent of payoff-relevant product characteristics, but rather by the internal structure

of the latter. This turns out to imply different equilibrium implications of default architecture. Consider the following variation on the model of this section. Each consumer is initially assigned to one of the two firms (each firm receives half the consumer population), as in the "opt out" default design described in Section 2. The consumer switches away to the other firm only when it offers a quality vector that dominates the one offered by his default firm.

This market model induces the *same payoff function* for the firms as case (*ii*) (except in the case of a tie along one dimension). As a result, symmetric Nash equilibrium is given by Proposition 6 - i.e., the same as if consumers cannot choose by default and choose according to a random single dimension. By assumption, there is full market participation under "opt out". However, since in equilibrium no market alternative ever dominates the other, consumers never switch away from their default options.

Alternatively, impose an "opt in" policy. That is, the consumer is initially assigned to an outside option that is equivalent to the quality vector $(0, 0)$, hence it is clearly inferior to market alternatives. However, when no market alternative dominates the other, this is a "hard choice" for the consumer - akin to incomparable formats in the model of Section 2 - to which the consumer responds by sticking to his (inferior) default option. In symmetric Nash equilibrium, firms offer $\bar{q} = 1$, with an arbitrary breakdown into the two quality dimensions. Since no market alternative dominates the other, the consumer adheres to his default option, and thus there is no market participation at all.

Suppose that we slightly perturb the "opt in" rule, such that a tiny fraction of the consumer population are initially assigned to one of the two firms (corresponding to small α in the model of Section 2). Now there is a symmetric Nash equilibrium in which firms play a mixed strategy that induces a true quality close to 1, but domination occurs with probability $\frac{1}{3}$ (Bachi and Spiegler (2014) conjecture that this is the highest domination rate that

is possible in equilibrium), such that the overall market participation rate in equilibrium is slightly above $\frac{1}{3}$. Thus, if we measure "true" consumer utility by the average quality of the alternative they end up with, then consumer surplus in this equilibrium is approximately $\frac{1}{3}$. In contrast, consumer surplus in the symmetric equilibrium under the "opt out" rule is $\frac{1}{2}$. In this sense, "opting out" outperforms "opting in", in contrast to the model of limited comparability analyzed in Section 2. Once again, we see that the equilibrium implications of default architecture are sensitive to the way we model the limited comparability problem that underlies consumer inertia.

5 Conclusion

Part of the appeal of "nudging" is that it seems to offer a "regulatory free lunch": helping consumers without infringing contractual freedom. This paper has demonstrated that at least theoretically, equilibrium market responses to "nudges" can eat away part of this free lunch, and potentially reverse the intended consequences. Moreover, the equilibrium analysis is sensitive to the procedural model underlying the very biases that nudging addresses. Accepting this critique means facing once again the stark dilemma between protecting boundedly rational consumers from market exploitation and maintaining contractual freedom.

At a certain level, the claim in this paper is obvious and familiar to economists from old debates about policy evaluation: when analyzing the theoretical consequences of an intervention, one should think about agents' equilibrium reaction to the intervention in terms of an explicit "structural" model that accounts for agents' observed stimulus-response patterns. However, the sense in which the paper is "structural" is quite unusual. Economists usually reserve the term for rational-choice models that are explicit about agents' preferences and information. The models in this paper are "structural" in the sense that they are based on choice procedures that in-

volve non-standard primitives, such as the ability to make comparisons or the tendency to procrastinate in the presence of hard choices.

Such a marriage between "behavioral" and "structural" approaches was promoted by Rubinstein (1998), albeit in a very different style and using very different terminology (and certainly without any intended policy relevance). However, it does not seem to be the norm among practitioners of behavioral economics, least of all in the field of law and economics. Hopefully, this paper demonstrated that adopting such an approach enriches the theoretical discussion of consumer protection in the presence of "behavioral" effects, even if it cannot by its very nature offer easy solutions.

References

[1] Armstrong, M. (2008), "Interactions between Competition and Consumer Policy," Competition Policy International 4, 97-147.

[2] Armstrong. M. and J. Vickers (2012), "Consumer Protection and Contingent Charges," Journal of Economic Literature, American Economic Association 50, 477-93.

[3] Bachi, B. and R. Spiegler (2014), "Buridanic Competition", mimeo, Tel Aviv University.

[4] Bar-Gill, O. (2012), Seduction by Contract: Law, Economics, and Psychology in Consumer Markets, Oxford University Press.

[5] Beshears, J., J. Choi, D. Laibson and B. Madrian (2012), "Simplification and saving," Journal of Economic Behavior & Organization 95, 130-145.

[6] Camerer, C., S. Issacharoff, G. Loewenstein, T. O'Donoghue and M. Rabin (2003), "Regulation for conservatives: behavioral economics and

the case for 'asymmetric paternalism'", University of Pennsylvania Law Review 151, 1211-1254.

[7] Chioveanu, I. and J. Zhou (2013), "Price Competition with Consumer Confusion," Management Science 59, 2450-2469.

[8] De Clippel, G., K. Eliaz and K. Rozen (2013), "Competing for Consumer Inattention," CEPR Discussion Papers 9553.

[9] DellaVigna, S. and U. Malmendier (2004), "Contract Design and Self-control: Theory and Evidence," The Quarterly Journal of Economics 119, 353-402.

[10] Ellison, G. (2006), "Bounded Rationality in Industrial Organization," in Blundell, Newey and Persson (eds.), Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press.

[11] Gabaix, X. and D. Laibson (2006), "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," The Quarterly Journal of Economics 121, 505-540.

[12] Grubb, M. (2012), "Consumer Inattention and Bill-Shock Regulation," MIT Sloan Research Paper No. 4987-12.

[13] Heidhues, P. and B. Koszegi (2010), "Exploiting Naivete about Self-Control in the Credit Market," American Economic Review 100, 2279-2303.

[14] Huck, S. and J. Zhou (2011), "Consumer Behavioural Biases in Competition – A Survey", Office of Fair Trade, UK.

[15] Kamenica, E., S. Mullainathan and R. Thaler (2011), "Helping Consumers Know Themselves," American Economic Review 101, 417-22.

[16] Iyengar, S., G. Huberman, and W. Jiang. (2004), "How Much Choice is Too Much: Determinants of Individual Contributions in 401K Retirement Plans." In Pension Design and Structure: New Lessons from Behavioral Finance, ed. O. S. Mitchell and S. Utkus, 83-95. Oxford: Oxford University Press.

[17] Lucas, R. (1976), "Econometric policy evaluation: A critique," Carnegie-Rochester Conference Series on Public Policy 1, 19-46.

[18] Mullainathan, S., J. Schwartzstein and W. Congdon (2012), "A Reduced-Form Approach to Behavioral Public Finance," Annual Review of Economics 4, 511-540.

[19] Muris, T. (2002), "The Interface of Competition and Consumer Protection," paper presented at Fordham Corporate Law Institute's 29th Annual Conference on International Antitrust Law and Policy, New York.

[20] Piccione, M. and R. Spiegler (2012), "Price Competition under Limited Comparability," *Quarterly Journal of Economics* 127, 97-135.

[21] Rubinstein. A. (1998), Modeling Bounded Rationality. MIT Press.

[22] Spiegler R. (2006), "Competition over Agents with Boundedly Rational Expectations," Theoretical Economics 1, 207-231.

[23] Spiegler, R. (2011), Bounded Rationality and Industrial Organization. Oxford University Press.

[24] Thaler, R. and C. Sunstein (2003), "Libertarian paternalism," American Economic Review 93, 175-179.

[25] Thaler, R and C. Sunstein (2008), Nudge: Improving Decisions about Health, Wealth, and Happiness, Yale University Press.

6 Appendix: Proofs

6.1 Proposition 1

Consider a symmetric Nash equilibrium strategy. Let F denote the marginal equilibrium distribution over prices. For any p in the support of F, define $\sigma_p(x)$ as the probability the equilibrium strategy assigns to the format x conditional on p. I begin with a few preliminary observations. First, note that since $\alpha > 0$, firms can secure a strictly positive profit by charging $p > 0$ and playing each format with positive probability. Therefore, $p = 0$ is not in the support of F. Second, F cannot have an atom on any $p > 0$ - otherwise, a firm can profitably deviate to a strategy that consists of a price $p' = p - \varepsilon$ and the mixed format strategy σ_p, where $\varepsilon > 0$ is arbitrarily small. Finally, the support of F must be an interval $[p_l, 1]$, where $p_l > 0$ - otherwise, if there is a "hole" (p, p') in the support of F, the strategy consisting of the price p' and the mixed format strategy σ_p generates a strictly higher payoff than (p, σ_p), which belongs to the support of the equilibrium strategy, a contradiction.

Let us now show that the overall probability that each format is played is $\frac{1}{2}$ - that is, $\int \sigma_p(a)dF(p) = \int \sigma_p(b)dF(p) = \frac{1}{2}$. Assume the contrary - i.e., w.l.o.g a is played with probability above $\frac{1}{2}$. Note that

$$\begin{aligned}
\sigma_1 &\in \arg\min_{\sigma} \int_{p_l}^{1} \sum_x \sum_y \sigma(x)\sigma_p(y)\pi^u(x,y)dF(p) \\
\sigma_{p_l} &\in \arg\max_{\sigma} \int_{p_l}^{1} \sum_x \sum_y \sigma(x)\sigma_p(y)\pi^u(x,y)dF(p)
\end{aligned}$$

Therefore, $\sigma_1(b) = \sigma_{p_l}(a) = 1$. Moreover, by continuity, $\sigma_p(b) = 1$ (0) for every p sufficiently close to 1 (p_l). Now consider the highest price $p \in [p_l, 1]$ for which $\sigma_p(a) > 0$. If a firm switches to the pure strategy (p, b) it necessarily increases its payoff: on one hand, if the rival firm's realized price is $p' < p$, then by assumption it is more likely to play a, such that playing b would

induce a lower comparison probability; and on the other hand, if the rival firm's realized price is $p' > p$, then by the definition of p, it plays b with probability one, such that playing b would raise the comparison probability from 0 to 1. In both scenarios, the deviation raises the firm's market share, hence it is profitable. It follows that both formats are played with probability $\frac{1}{2}$.

Our next step is to show that $\sigma_p(a) = \frac{1}{2}$ for almost every $p \in [p_l, 1]$. Assume the contrary - i.e., w.l.o.g there is a price $p \in [p_l, 1]$ for which

$$\int_p^1 \sigma_p(a) dF(p' \mid p' > p) > \frac{1}{2}$$

Then, it must be the case that there is such p satisfying $\sigma_p(a) > 0$. Since the overall probability that a is played is exactly $\frac{1}{2}$, it must be the case that

$$\int_p^1 \sigma_p(a) dF(p' \mid p' < p) < \frac{1}{2}$$

Therefore, if a firm deviates to the pure strategy (p, b), the deviation raises (lowers) comparison probability against higher (lower) price realizations of the rival firm, hence it is profitable.

The result that $\sigma_p(a) = \frac{1}{2}$ for almost all p implies that for almost any profile of realized prices, the comparison probability is $\frac{1}{2}$. This enables us to determine F. When a firm charges $p = 1$, its payoff is

$$1 \cdot \frac{\alpha}{2} \cdot \frac{1}{2} = \frac{\alpha}{4}$$

Thus, for every $p \in [p_l, 1]$, a firm's payoff should satisfy

$$p \cdot \left[\frac{\alpha}{2} \left(\frac{1}{2} + \frac{1}{2}(1 - F(p)) \right) + \left(1 - \frac{\alpha}{2} \right) \frac{1}{2}(1 - F(p)) \right] = \frac{\alpha}{4}$$

This equation gives us F, and p_l is derived from the equation $F(p_l) = 0$.

6.2 Proposition 2

First, observe that the preliminary observations made at the beginning of the proof of Proposition 1 hold, with one exception: we cannot rule out the possibility that F places an atom on p and $\sigma_p(a) = 1$. The reason is that since $\pi^c(a,a) = 0$, deviating to the pure strategy $(p - \varepsilon, a)$ is not profitable for an arbitrarily small $\varepsilon > 0$.

Let $p_h(x)$ and $p_l(x)$ denote the highest and lowest prices p in the closure of the set $\{p \in [p_l, 1] \mid \sigma_p(x) > 0\}$. Let us show that $p_h(b) \leq p_l(a)$. Assume the contrary, namely that there exist $p^1, p^2 \in [p_l, 1]$ such that $p^2 > p^1$, $\sigma_{p^2}(b) > 0$ and $\sigma_{p^1}(a) > 0$. For each $k = 1, 2$, $x = a, b$, the market share that the strategy (p^k, x) generates, denoted $s(p^k, x)$, is as follows (to simplify the notation, I ignore the possibility of an atom on p^k - to incorporate atoms we would have to replace F with left or right limits of F - without changing the argument):

$$\begin{aligned} s(p^k, x) &= \frac{\alpha}{2}\left[1 - \int_{p_l}^{p^k}\left(\sum_y \sigma_p(y)\pi^c(x,y)\right)dF(p)\right] \\ &\quad + \left(1 - \frac{\alpha}{2}\right)\int_{p^k}^{1}\left(\sum_y \sigma_p(y)\pi^c(x,y)\right)dF(p) \end{aligned}$$

Since $\pi^c(b,x) > \pi^c(a,x)$ for every x, it is straightforward to verify that it is impossible that $s(p^1, a) \geq s(p^1, b)$ and $s(p^2, b) \geq s(p^2, a)$, because F assigns positive probability to the interval (p^1, p^2).

Suppose that $p_l(a) < 1$. We have just seen that $\sigma_p(a) = 1$ for every $p \in (p_l(a), 1)$. Since $\pi^c(a,a) = 0$, if a firm deviates from a price $p \in (p_l(a), 1)$ to the pure strategy $(1, a)$, its market share does not change and hence its payoff increases. Therefore, $p_l(a) = 1$, such that $\sigma_p(b) = 1$ for almost every $p \in [p_l, 1)$. If F does not place an atom on $p = 1$, this means that b is played with probability one, and in this case a firm that charges a price close to 1 will strictly prefer to deviate to a. Thus, it must be the case that F places an atom on $p = 1$. To calculate the size of this atom λ, observe that the

following equality must hold:

$$s(1,a) = \frac{\alpha}{2} \cdot \left[\lambda + \frac{1}{2}(1-\lambda)\right] = \frac{\alpha}{2} \cdot \lambda + \left(1 - \frac{\alpha}{2}\right) \cdot \frac{1}{2}\lambda = \lim_{\varepsilon \to 0} s(1-\varepsilon, b)$$

Otherwise, there would be a profitable deviation either from $(1,a)$ to $(1,b)$, or from (p,b) to (p,a) for some p sufficiently close to 1. Thus, $\lambda = \frac{\alpha}{2}$, such that a firm's payoff from $(1,b)$ is $\frac{\alpha}{2}(\frac{1}{2} + \frac{\alpha}{4})$. Since this is the equilibrium payoff, it is the payoff from (p,a) for every $p \in [p_l, 1)$. Thus, we can write:

$$p \cdot \left[\frac{\alpha}{2}(1 - F(p)) + \left(1 - \frac{\alpha}{2}\right)\left(\frac{1}{2}\lambda + (1 - \lambda - F(p)\right)\right] = \frac{\alpha}{2}(\frac{1}{2} + \frac{\alpha}{4})$$

and retrieve the expression for F, as well as the value of p_l, from this equation.

6.3 Proposition 3

As DellaVigna and Malmendier (2004) show, in this environment the symmetric equilibrium tariff (A^*, p^*) maximizes consumers' perceived ex-ante net utility

$$u(x^u) - A - px^u$$

subject to the zero-profit condition

$$A + (p - \frac{1}{4})x^v = 0$$

where $x^u = \arg\max_x u(x) - px$ and $x^v = \arg\max_x v(x) - px$. Since v is concave and twice differentiable, x^v is simply given by the first-order condition $v'(x) = p$, hence $x^v = \frac{1}{p} - 1$. In contrast, $x^u = x^v < \varepsilon$ for $p > \frac{1}{1+\varepsilon}$, and $x^u = \varepsilon$ for $\frac{1}{1+\varepsilon}$. If $p^* > \frac{1}{1+\varepsilon}$, then the constrained maximization problem is reduced to maximizing $u(x) - \frac{1}{4}x$, and it is easy to see that p^* cannot be an optimum because it is outperformed by a slightly lower p. If $p^* \leq \frac{1}{1+\varepsilon}$, the constrained maximization problem is reduced to minimizing A subject to the zero-profit condition. This means that p^* should maximize $(p - \frac{1}{4})(\frac{1}{p} - 1)$, and we obtain

$p^* = \frac{1}{2}$.

6.4 Proposition 4

Let $t^*(x) = \frac{1}{4}x$ for all x. Let us first show that $(t^*, 3)$ is stable. Suppose that a firm deviates to some other t defined by (A, p). In order for consumers to choose t' over t^*, it must be the case that

$$A + 3p < \frac{3}{4}$$

Given p, consumers who select t' will subsequently choose $x = \arg\max[\ln(1+x) - px]$, hence $x = \frac{1}{p} - 1$. In order for the deviation to be profitable, we must have

$$A + (p - \frac{1}{4})(\frac{1}{p} - 1) > 0$$

Combining the two inequalities, we obtain

$$(\frac{1}{4} - p)(\frac{1}{p} - 1) < 3(\frac{1}{4} - p)$$

and there is no p that satisfies this condition.

Let us now show that there is no other stable pair (t^*, x^*), where t^* is defined by (A^*, p^*), and x^* is given by $v'(x^*) = p^*$, namely $x^* = \frac{1}{p^*} - 1$. First, let us show that a necessary condition for stability is that firms earn zero profits - i.e., $A^* + (p^* - \frac{1}{4})x^* = 0$. If $A^* + (p^* - \frac{1}{4})x^* < 0$, a firm can deviate to (A, p^*) where $A > A^*$; consumers will not choose the firm, and so it will make zero profits, hence the deviation is profitable. If $A^* + (p^* - \frac{1}{4})x^* > 0$, a firm can deviate to $(A^* - \varepsilon, p^*)$, where $\varepsilon > 0$ is arbitrarily small; consumers will select the firm because it obviously has a lower effective price, and they will proceed to choose x^* because the price-per-unit is still p^*, hence the deviation is profitable.

The pair (t^*, x^*) is unstable if and only if we can find (A, p, z^*), such that

$$\begin{aligned} A + px^* &< A^* + p^*x^* \\ A + (p - \frac{1}{4})z^* &> A^* + (p^* - \frac{1}{4})x^* \end{aligned}$$

where $z^* = \frac{1}{p} - 1$. Since $A^* + (p^* - \frac{1}{4})x^* = 0$, the inequalities can be simplified into

$$(p - \frac{1}{4})(\frac{1}{p} - 1) > (p - \frac{1}{4})(\frac{1}{p^*} - 1)$$

By zero If $p^* > \frac{1}{4}$, any $p \in (\frac{1}{4}, p^*)$ satisfies it, whereas if $p^* < \frac{1}{4}$, any $p \in (p^*, \frac{1}{4})$ satisfies it. This completes the proof.