On the Equilibrium Effects of Nudging

Ran Spiegler

ABSTRACT

Consumers' systematic decision biases make them vulnerable to market exploitation. The doctrine of libertarian paternalism maintains that this problem can be mitigated by soft interventions (nudges) 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 notion of bounded rationality. A reasonable consumer making decisions under normal circumstances is not infallible: he is naturally limited in his computational abilities; his understanding of market regularities is imperfect; he often suffers from attention deficits because of the many tasks he needs to juggle; and he may succumb to temptation, self-delusion, or wishful thinking. Since these limitations make him vulnerable 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 (such as sin taxes and banning contracts ex ante or voiding them ex post)?

One school of thought is encapsulated in the words of a former Federal Trade Commission chairman: "Robust competition is the best single

RAN SPIEGLER is a professor at Tel Aviv University, University College London, and the Centre for Macroeconomics. This paper has enjoyed financial support from the Pinhas Sapir Center for Development. It grew out of a project that was supported by European Research Council grant no. 230251. I thank Ariel Ben-Porat, Oren Bar-Gill, Kfir Eliaz, Michael Grubb, Ariel Rubinstein, seminar participants at Tel Aviv University and Yale University, and the editor and referee of this journal for their useful comments.

[*Journal of Legal Studies*, vol. 44 (June 2015)]
 0047-2530/2015/4402-0024$10.00

means for protecting consumer interests" (Muris 2002). However, the growing literature on behavioral industrial organization (BIO) does not provide unambiguous theoretical support for this sweeping motto, for a variety of reasons.

First, when consumers are diversely sophisticated, stronger competition may impel firms to shift the focus of their pricing strategy toward exploitation of naive consumers (Varian 1980; Spiegler 2011, ch. 4.2; Armstrong 2015). Second, in a market for a product of questionable intrinsic value (certain alternative-medicine practices and active money management are often claimed to fall into this category), increased supply may end up increasing the fraction of consumers who take the welfare-reducing decision to enter this market (Spiegler 2006b). Third, firms may respond to stronger competition by intensifying their obfuscation tactics rather than by offering more attractive products (Spiegler 2006a; Carlin 2009; Chioveanu and Zhou 2013).

In the past decade, a new school of thought (Thaler and Sunstein 2003; Camerer et al. 2003; Bar-Gill 2012) has argued for a third way that consists of soft interventions that assist boundedly rational consumers without constraining contractual freedom. According to this approach, 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 or the specification of a default option. The regulator could also impose user-friendly disclosure requirements that reduce information-processing effort and minimize confusion. Choice architecture and disclosure are indeed the prime examples of this approach, dubbed "libertarian paternalism" or "nudging," which received a major boost by Thaler and Sunstein's eponymous best seller (Thaler and Sunstein 2008).

The "nudgniks" have been criticized on philosophical grounds. In particular, it has been argued that any governmental manipulation of consumer choice goes against libertarian values, especially when it springs from a paternalistic pretense to know consumers' true preferences.[1] This paper offers a different critique, which targets two characteristics of most discussions of nudging. First, although proponents of nudging recognize that beneficial effects of nudges may be counteracted by market equilibrium responses (Barr, Mullainathan, and Shafir 2008; Bar-Gill 2012, pp. 38–39, 175–76; Baker and Siegelman 2014), equilibrium analyses of

1. Thaler and Sunstein (2008) and Sunstein (2014) discuss and reply to these criticisms.

nudges are rare. Second, the literature typically regards consumers' decision errors and biases as primitive behavioral phenomena or black boxes. For example, Bar-Gill (2012) and Mullainathan, Schwartzstein, and Congdon (2012) employ models in which rational choice involves trading off costs and benefits according to certain decision weights, and decision biases are captured by wrong weights. Although Mullainathan, Schwartzstein, and Congdon 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 the origin of the wrong weights, thus offering little guidance as to how they could be affected by firms' obfuscation tactics or regulatory interventions. The two characteristics are thus interrelated: the reduced-form approach to modeling consumer biases limits the scope of equilibrium analysis of nudging.

I challenge this viewpoint by considering a sequence of market models in which profit-maximizing firms compete for boundedly rational consumers. In each model, consumers commit a decision error that intuitively calls for nudging. However, the model does not treat the error as a black box but formalizes a plausible, explicit psychological mechanism that generates it.[2] I use each model to address a different nudge, which seems pertinent given the decision error in question. In each case, I show that once the psychological mechanism behind consumer error and the firms' equilibrium response are taken into account, the theoretical case for the nudge is reversed. I now preview this sequence of exercises.

1.1. Default Architecture and Limited Comparability

Decision makers often adhere to status quo or default options, even when there are better alternatives and when the physical cost of switching away from the default is negligible. Such default bias leads to social outcomes such as low consumer switching rates in retail banking or weak participation in retirement saving and organ donation programs. Consequently, design of default rules has been a major theme in the literature on nudging. To evaluate the equilibrium implications of default architecture in consumer-market settings, we should look for the deeper psychological mechanism behind default bias. Introspection and experimental evidence suggest that complexity of the choice problem is an important contributing factor to default bias. In a market environment, this complexity is affected by the firms' obfuscation strategies. Adapting a modeling ap-

2. Rubinstein (1998) refers to such choice models as "procedural."

proach that originates in the classic paper Varian (1980) and is developed in Carlin (2009), Piccione and Spiegler (2012), and Chioveanu and Zhou (2013), I analyze a two-firm price-competition game in which firms also choose whether to obfuscate, where obfuscation is simply an action that reduces the probability that consumers are able to make a price comparison (thus increasing the probability that they will adhere to their default option, whatever it is). I show that default rules aimed at increasing market participation may end up reducing consumer welfare, as a result of firms' equilibrium price and obfuscation responses.

1.2. Product-Use Disclosure and Underestimation of Future Taste for Immediate Gratification

A common concern is that when consumers accept long-term service contracts (such as credit cards or mobile phone plans) they systematically underestimate the total amount they will end up paying. It has been argued (Thaler and Sunstein 2008; Bar-Gill 2012) that regulators can respond by mandating produce-use disclosure. For instance, firms can be required to report the effective price per unit induced by the tariffs they offer, calculated according to average consumption in a comparison group of consumers. To evaluate this nudge, I follow the suggestion in Bar-Gill (2012) that excessive payments may originate from consumers' underestimation of their own future taste for immediate gratification, as captured in the influential model of DellaVigna and Malmendier (2004). I show that when this model is extended to incorporate product-use disclosure, consumers' excessive payment is exacerbated and their ex ante welfare drops. The reason is that the disclosure effectively impels firms to cater to consumers' taste for immediate gratification rather than to their ex ante preferences.

1.3. Product-Attribute Disclosure and Trade-off Avoidance

Consumers often appear to neglect nonsalient product attributes (certain bank fees, contingencies in an insurance contract, and add-ons). Firms can take advantage of such neglect and make shrouded attributes unattractive for consumers. Moreover, Gabaix and Laibson (2006) argue that market forces need not impel firms to unshroud these attributes. A natural regulatory response is to mandate full disclosure of all product attributes to homogenize their salience. I argue that if the deeper psychological force behind consumers' neglect of attributes is an intrinsic attention deficit or aversion to hard trade-offs, disclosure will not turn consumers into rational trade-off machines; instead, they will continue to neglect attributes, albeit less predictably because of the attributes' uniform salience. This lack of predictability weakens competitive pressures and can make consumers worse off. I articulate this point using the model in Bachi and Spiegler (2014).

I wish to make a few comments regarding the methodology of this paper. First, the models I have selected are, in their basic form, among the most well known in the BIO literature. They are extended to accommodate the nudges, but I have striven to make the extension as minor as possible, and I have chosen the models' simplest possible form to illustrate the nudges' effects, sacrificing generality for the sake of expositional clarity and simplicity. Thus, one should not expect the theoretical exercises in this paper to be directly applicable to concrete policy issues. Second, I have deliberately chosen to use different models involving different psychological mechanisms to examine different nudges. This will hopefully convince the reader that the anomalous equilibrium effects of the nudges are not cooked but are plausible consequences of taking equilibrium analysis seriously. Some of the models can be used to address multiple nudges, as illustrated in the online appendix. Finally, the scope of my analysis is limited in the sense that it focuses on consumer markets. However, I believe that it is relevant to other settings as well. For example, the equilibrium analysis of default architecture sheds light on the role of employers in mediating the interaction between retirement saving funds and savers: in their absence, default architecture could backfire because of equilibrium effects.

One advantage of thinking about consumers' decision errors in terms of their deeper procedural origins is that it broadens the range of consumer-protection measures that can be studied, and it enables us to make finer distinctions between regulatory environments. In the online appendix, I discuss further implications of the modeling approaches in Sections 3 and 5 for various nudges.

2. RELATED LITERATURE

The BIO literature forms the backdrop for this paper; see Ellison (2006), Armstrong (2008), and Huck and Zhou (2011) for surveys and Spiegler (2011) for a graduate-level textbook. The latter provides many theoretical examples of regulatory interventions that seem at first glance to promote consumer welfare (strengthening competition or requiring firms to offer plain-vanilla options) and yet end up harming consumers when their underlying choice procedures and the firms' equilibrium response are taken into account (see Sections 4 and 7, for example). However, since these interventions are not libertarian-paternalistic, I do not discuss them in detail in this paper.[3]

A number of recent works examine the equilibrium effects of regulated disclosure in the contexts of market models in which (some) consumers have various forms of limited attention or awareness. Kamenica, Mullainathan, and Thaler (2011) show that product-use disclosure may have an adverse welfare effect on consumers with limited knowledge of their own tastes, once market equilibrium effects are taken into account. Armstrong and Vickers (2012) discuss disclosure of contingent charges in a model in which some consumers neglect nonsalient contingencies. De Clippel, Eliaz, and Rozen (2013) analyze a model in which consumers rationally allocate their limited attention to individual markets according to the observed prices charged by market leaders. They showed 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. Grubb (2015) constructs a dynamic consumption model in which consumers have limited ability to monitor their own past consumption and analyzes disclosure policies that address the bill-shock problem that arises in such a model. In particular, he shows that policies that prevent bill shock can reduce consumer welfare by effectively robbing firms of instruments for discriminating among consumers with different demand.

What distinguishes the present paper from these works? First, it presents the equilibrium critique of nudging in a comprehensive and systematic manner. Second, it contains the first theoretical analysis of equilibrium effects of default architecture (together with Bachi and Spiegler [2014]). Finally, most interventions studied in the BIO literature are relevant in models with rational consumers. In contrast, nudges are by definition irrelevant when the consumer is rational. The nudges studied in Sections 3 and 5 can be viewed precisely as external manipulations of the very parameters that define the consumer's bounded rationality.

Another recent development is the empirical investigation of the equilibrium effects of nudging. Duarte and Hastings (2012) study the impli-

3. Heidhues and Koszegi (2010) analyze the welfare implications of a "hard" paternalistic intervention, banning late fees, in credit markets when consumers mispredict their future self-control.

cations of disclosure regulations on the structure of management fees charged by privatized social security funds in Mexico. Grubb and Osborne (2015) estimate a model of the mobile phone market based on Grubb (2015) and use it to assess the welfare effects of product-use regulation in that industry. Handel (2013) estimates a model of health insurance choice with an explicit adverse-selection component and a reduced-form account of consumer inertia and uses the estimation to argue that nudges can backfire because they exacerbate the adverse-selection problem.

3. 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 or status quo options (even when the physical cost of switching is negligible) in a variety of contexts: donating organs, saving for retirement, renewing insurance policies, or selecting a provider of retail banking services. This has given rise to the suggestion, eloquently articulated by Thaler and Sunstein (2008), that policy makers can and should influence decision makers' choices in such settings by appropriately designing default options. In particular, it has been argued that a switch from opting in (a regime in which the default is an outside option) to opting out (a regime in which the default is one of the market alternatives) will increase participation rates in various programs. Thaler and Sunstein (2008) acknowledge that default bias is not a primitive phenomenon and that it originates from loss aversion, limited attention, or choice complexity. However, neither they nor others integrate these considerations explicitly into an equilibrium analysis of default design.

In this section I attempt such an exercise in the context of a market model in which profit-maximizing firms compete for consumers who exhibit default bias, and I focus on choice complexity as the source of this bias. This is an extension of the classic model of Varian (1980), in which firms selling a homogenous product play a price-competition game. Unlike the textbook model of Bertrand competition, not all consumers select the cheapest market alternative. Instead, a certain fraction α of the consumer population randomly chooses one of the firms, independently of their prices. Varian interprets this behavior as a manifestation of limited information or customer loyalty. More recently, Carlin (2009), Piccione and Spiegler (2012), and Chioveanu and Zhou (2013) reinterpret Varian's friction in terms of limited ability to make price comparisons and extend the model by allowing α to depend on the firms' endogenous marketing strategies. I extend this tradition by incorporating default architecture.

3.1. The Model

Consider a market that consists of two profit-maximizing firms and a measure one of consumers. The firms costlessly produce a homogenous product that has a gross value of 1 for consumers, who also have an outside option that gives them 0 net utility. The firms play a simultaneous-move game with complete information. Each firm $i = 1, 2$ chooses a pair (p_i, x_i), where $p_i \in [0, 1]$ is the price of the firm's product and $x_i \in \{0, 1\}$ indicates firm i's obfuscation strategy, such that $x_i = 1$ (0) means that the firm obfuscates (refrains from obfuscation).

The effect of playing $x = 1$ is to lower the probability that consumers are able to make a price comparison between the two market alternatives. Specifically, each consumer makes a comparison with probability $1 - \frac{1}{2}(x_1 + x_2)$. This particular formula is assumed purely for expositional simplicity. It means that when both firms refrain from obfuscation, comparability is perfect, as in the Bertrand model; when both firms obfuscate, the consumer is totally unable to make a comparison; and when only one firm obfuscates, the consumer makes a comparison with probability $\frac{1}{2}$. Whenever a consumer cannot make a comparison, I say that he faces a complex choice.

This formulation approximates a variety of real-life forms of obfuscation. For instance, contractual terms can be described in technical or jargon-laden language that requires translation, and the act of translation can result in gibberish that prevents comparison. Likewise, prices can be presented in way that requires the consumer to perform a complex calculation to derive the bottom line, and the calculation can break down.

I allow for heterogeneity among consumers in terms of their response to complex choices. A fraction $\lambda \in (0, 1]$ of consumers are decisive types—they arbitrarily choose a firm (each with probability $\frac{1}{2}$) whenever they cannot make a comparison; the remaining fraction $1 - \lambda$ are indecisive types—they respond to complex choices by deciding not to decide, namely, choosing a default option when possible (and when choosing by default is infeasible, they are forced to choose decisively). A lower value of λ means that the propensity for default bias is more common in the consumer population, and $\lambda = 1$ means that consumers exhibit no default bias. The default regime is designed ex ante by the regulator. I focus on two default rules. Under opt in, the default is the outside option. Under opt out, the default is firm 1 for half the consumer population and firm 2 for the other half (to maintain the game's symmetry, for simplicity).[4]

I now define the payoff function in the simultaneous-move game the firms play. Firm i's profit is $p_i \times s_i\,[(p_1, x_1), (p_2, x_2)]$, where s_i represents the firm's market share. Under opt in, s_i is given by

$$s_i[(p_1,\ x_1),\ (p_2,\ x_2)] = \begin{cases} \dfrac{x_1 + x_2}{2} \times \dfrac{\lambda}{2} + \left(1 - \dfrac{x_1 + x_2}{2}\right) \times 1 & \text{if } p_i < p_j \\ \dfrac{x_1 + x_2}{2} \times \dfrac{\lambda}{2} + \left(1 - \dfrac{x_1 + x_2}{2}\right) \times \dfrac{1}{2} & \text{if } p_i = p_j \\ \dfrac{x_1 + x_2}{2} \times \dfrac{\lambda}{2} + \left(1 - \dfrac{x_1 + x_2}{2}\right) \times 0 & \text{if } p_i > p_j. \end{cases} \quad (1)$$

Under opt out, each firm ends up having half the consumers who are unable to make a comparison (some of them are decisive and choose the firm arbitrarily, while others are indecisive and were initially assigned to the firm as a default). Consequently, the market share function s_i is also given by equation (1), except that I substitute $\lambda = 1$. Thus, as far as firms are concerned, opt out is equivalent to opt in with $\lambda = 1$. For this reason, I identify opt out with $\lambda = 1$ when I conduct the equilibrium analysis and return to the distinction between the two default regimes in the welfare analysis.

What is the relation between this market model and the various contexts in which default architecture has been discussed? First, nonmarket activities such as organ donation are clearly outside the model's scope. In markets for long-term services (insurance, magazine subscriptions, or mobile phone services), opt in may correspond to a regulatory intervention that rules out automatic contract renewals, whereas opt out fits an environment in which autorenewals are the norm. With regard to 401(k) retirement saving programs, an important qualification is in order. In reality, funds do not compete directly for savers; instead, the interaction is mediated by the savers' employers, which shape the set of feasible alternatives and its presentation and negotiate the management fees. This market is thus effectively regulated by the employers. The analysis in this section can be viewed as an attempt to speculate about the equilibrium effects of default architecture in the absence of such de facto regulation.

4. The case of $\lambda = 0$ results in multiple Nash equilibria under opt in, and therefore I rule it out.

3.1.1. Comment: Buridanic Behavior. The net value that firm i generates for consumers is $1 - p_i \geq 0$. Hence, by assumption, the outside option is inferior to any of the market alternatives, regardless of the firms' behavior. This is a deliberate modeling choice that is designed to ensure that market participation is unambiguously beneficial for consumers. It also means that under opt in, an indecisive consumer (who clings to the outside option in response to complex choices) ends up being worse off than if he chose an arbitrary market alternative. An indecisive consumer is like the proverbial Buridan's ass: unable to rank two attractive market options, he procrastinates choice, even though he may recognize that the delay is suboptimal because any of the market products is better than opting out. As a result, for a given time horizon, the effective choice of some of the consumers is not to enter. The psychology behind such behavior is that people generally dislike making an active choice that lacks a good reason (Payne, Bettman, and Johnson 1993; Anderson 2003), and they are willing to delay choice in order to avoid this unpleasant feeling. I believe that this behavior is quite common. For hard empirical evidence of Buridanic behavior in various contexts, see Iyengar, Huberman, and Jiang (2004), Madrian and Shea (2001), and Beshears et al. (2012).[5]

3.1.2. Comment: Active Choice. A third default rule discussed in the nudge literature is active choice, namely, forcing consumers to make an explicit choice and forbidding them to choose by default. In this case, given that the outside option is clearly inferior to any of the market alternatives, it makes sense to assume that when the consumer faces a complex choice, he behaves as if he were decisive. Under this assumption, active choice is payoff equivalent to the case of $\lambda = 1$ under any of the other default rules.

3.1.3. Comment: The Upper Bound on Prices. I have assumed that the upper bound on the firms' prices coincides with 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 overpricing. Another answer is

5. Because the outside option is inferior to the market alternatives, one could argue that the fraction of indecisive types should be lower under opt in than under opt out. However, this criticism is irrelevant for my purposes, because under opt out, the fraction of indecisive consumers does not matter for firms' payoffs and consumer welfare.

that consumers do know the utility from their default option, and since they switch only when having a good understanding of the choice problem, they will never end up paying more than they are willing to pay.

3.2. Symmetric Nash Equilibria

The game between the two firms has no pure-strategy equilibrium, regardless of the default rule. To see why, suppose first that both firms play the same (p, x). If $p = 0$, any firm can deviate to $p' > 0$, $x = 1$. As a result, a positive fraction of consumers will be unable to make a price comparison, and because $\lambda > 0$, the firm will necessarily have a positive clientele; hence, the deviation is profitable. If $p > 0$, any firm can slightly undercut the price and play $x = 0$. This will necessarily increase the firm's market share by a margin that is bounded away from 0 because the consumer will be able to make a comparison with probability $\frac{1}{2}$ at least; hence, the deviation is profitable.

I now turn to analyzing symmetric mixed-strategy Nash equilibria. A mixed strategy in this model is a joint probability distribution over pairs $(p, x) \in [0, 1] \times \{0, 1\}$. Every (p, x) in the support of a symmetric equilibrium strategy maximizes a firm's payoff, given that the other firm plays the same mixed strategy. It turns out that there is a unique symmetric equilibrium; the equilibrium strategy exhibits correlation between the firm's pricing and obfuscation decisions. With probability $\lambda/2$, it charges the monopoly price $p = 1$ and obfuscates $(x = 1)$. With the remaining probability $1 - \lambda/2$, the firm randomizes continuously over a range of prices $p < 1$ and refrains from obfuscation $(x = 0)$.

Proposition 1. The game has a unique symmetric Nash equilibrium. With probability $\lambda/2$, firms play $(p, x) = (1, 1)$. With the remaining probability $1 - \lambda/2$, they play $x = 0$ and mix over prices according to the cumulative distribution function

$$F(p) = M + 1 - \frac{M}{p},$$

defined over the interval $[M/(M + 1), 1]$, where

$$M = \frac{\lambda(2 + \lambda)}{4(2 - \lambda)}.$$

The intuition behind the firms' equilibrium behavior is as follows. When a firm chooses to charge a high price, it has an incentive to lower the probability of price comparison, and therefore it prefers to obfuscate. In contrast, when a firm chooses a low price, it seeks comparison and therefore refrains from obfuscating.[6] An increase in λ leads to less-competitive equilibrium behavior, including a higher probability of playing $p = 1$ and a higher $M/(M + 1)$. The reason is that when λ is high, consumers who face a complex choice typically end up choosing one of the firms, independently of the price profile. This means that firms benefit from choice complexity, so they have a big incentive to generate it by obfuscating and to exploit it by charging a high price.

3.3. Welfare Analysis

Recall that the outside option generates 0 net utility for consumers, while the market alternatives are produced costlessly and generate a gross utility of 1 for consumers. Therefore, social surplus is simply equal to the market participation rate, namely, the total fraction of consumers who end up choosing a market alternative. Under opt out, consumers always end up choosing one of the firms; hence, the equilibrium participation rate is 1. To calculate the equilibrium participation rate under opt in, note first that the fraction of consumers who make a price comparison in the symmetric Nash equilibrium is

$$\left(1-\frac{\lambda}{2}\right)^2 \times 1 + 2 \times \frac{\lambda}{2}\left(1-\frac{\lambda}{2}\right) \times \frac{1}{2} + \left(\frac{\lambda}{2}\right)^2 \times 0 = 1 - \frac{\lambda}{2}.$$

Thus, under opt in, the equilibrium market participation rate in equilibrium is

$$\lambda + (1-\lambda)\left(1-\frac{\lambda}{2}\right) = \frac{1}{2}\lambda^2 - \frac{1}{2}\lambda + 1. \tag{2}$$

The reason is that a fraction λ of the consumers end up with one of the market alternatives, regardless of their comparability, whereas the remaining fraction $1 - \lambda$ of the consumer population enters only when able to make a comparison. Expression (2) is U-shaped with regard to λ: it attains the maximum of 1 both at $\lambda = 1$ and in the $\lambda \to 0$ limit (and it attains a minimum of $\frac{7}{8}$ at $\lambda = \frac{1}{2}$). Thus, equilibrium social welfare is nonmonotone in consumers' propensity for default bias; in particular,

6. The cutoff price that separates obfuscatory and nonobfuscatory behavior is $p = 1$, and the equilibrium price distribution has an atom on this price. This is a fragile property because of the extreme assumption of 0 comparability when $x_1 = x_2 = 1$. If I perturb the model and assume a small positive comparison probability in this case, the equilibrium price distribution will have no discontinuous jumps; firms will play $x = 1$ over a small range of prices near $p = 1$, and the overall probability of $x = 1$ will be close to $\lambda/2$.

when virtually all consumers are indecisive, opt in and opt out induce the same social welfare in equilibrium.

The fact that there is full equilibrium participation when $\lambda = 1$ is not surprising, because this is built into the definition of the consumer's choice procedure. The more noteworthy observation is that the equilibrium participation rate converges to 1 in the $\lambda \to 0$ limit. Thus, if the regulator introduces the opt-in default regime and almost all consumers are indecisive, they end up exhibiting no default bias in equilibrium. The reason is that under opt in, firms do not benefit at all from default bias when $\lambda \to 0$, and therefore they have no incentive to obfuscate. As a result, consumers never face complex choices. This observation illustrates the importance of exploring equilibrium effects and the psychological mechanism behind decision errors. When the deep structural parameter λ reflects an extreme propensity for default bias (that is, λ is close to 0), equilibrium effects under opt in produce an observed behavior that exhibits virtually no default bias.

I now calculate equilibrium industry profits. Suppose that firm 1 considers playing the pure strategy $(p, x) = (1, 1)$. In equilibrium, firm 2 plays the same pure strategy with probability $\lambda/2$. In this case, $x_1 = x_2 = 1$, such that consumers are entirely unable to make a comparison, and each firm gets a fraction $\lambda/2$ of the consumer population. With probability $1 - \lambda/2$, firm 2 plays $p < 1$ and $x = 0$. Since comparison probability in this case is $\frac{1}{2}$, firm 1 gets a fraction of $\lambda/2 \times \frac{1}{2}$ of the consumer population. Thus, when firm 2 plays the equilibrium strategy, the firm's profit from the pure strategy $(p, x) = (1, 1)$ is

$$1 \times \left[\frac{\lambda}{2} \times \frac{\lambda}{2} + \left(1 - \frac{\lambda}{2}\right) \times \frac{\lambda}{2} \times \frac{1}{2}\right] = \frac{1}{8}\lambda^2 + \frac{1}{4}\lambda. \tag{3}$$

This pure strategy belongs to the support of the symmetric equilibrium mixed strategy. A basic property of mixed-strategy Nash equilibrium is that every pure strategy in the support of the equilibrium strategy is a best reply. Therefore, each firm earns $\frac{1}{8}\lambda^2 + \frac{1}{4}\lambda$ in equilibrium.

Net consumer welfare in equilibrium is equal to social surplus minus industry profits:

$$\left(\frac{1}{2}\lambda^2 - \frac{1}{2}\lambda + 1\right) - 2 \times \left(\frac{1}{8}\lambda^2 + \frac{1}{4}\lambda\right) = 1 + \frac{1}{4}\lambda^2 - \lambda.$$

This expression is decreasing in λ. It follows that opt in is superior to opt out in terms of equilibrium consumer welfare. In particular, in the $\lambda \to 0$ limit (that is, when the default regime is opt in and consumers have an extreme propensity for default bias), consumer welfare reaches the maximal level of 1, because market participation is full and equilibrium prices converge to 0.

The intuition for this result is that, in this model, default bias results from the complexity of price comparison. The opt-in rule restrains firms' incentive to obfuscate, because they benefit less from choice complexity (compared with opt out). In the limit case in which all consumers are indecisive, firms derive no benefit from choice complexity under opt in and refrain entirely from obfuscation. This in turn means that price competition is as transparent as in the Bertrand benchmark, such that equilibrium prices reach the competitive level. Thus, paradoxically, once we take the psychological mechanism behind default bias and the firms' incentives into account, the default rule that seems less conducive to market participation turns out to maximize both participation and net consumer welfare, precisely when consumers' propensity for default bias is at its extreme.

3.4. Discussion

The exercise in Section 3 demonstrates the importance of accounting for the psychological origins of consumers' decision errors. Default bias is not a primitive phenomenon: it is determined at least in part by consumers' attitudes to complex choice problems. I focus on the effect of limited comparability on default bias and show that a change in consumers' default specification affects firms' incentive to manipulate comparability, which in turn has an effect on competitive pressures.

What are the possible implications of this analysis for some of the real-life examples that motivate my discussion in this section? Recall the interpretation of default rules in terms of automatic renewal of long-term services such as insurance or magazine subscriptions. The equilibrium analysis provides support for the intuition that in these environments, banning autorenewals leads to higher levels of consumer welfare. It suggests that in response to such a regulation, firms have a weaker incentive to obfuscate, and this in turn can lead to high market participation and low prices.

With respect to retirement savings programs such as 401(k) plans, my exercise highlights the importance of the employer's role in mediating the interaction between funds and savers. It suggests that the reason that a switch from opt in to opt out may increase participation rates without harming savers is that the employer acts as a de facto market regulator. Savers benefit not only from the soft paternalism of default architecture but also from the employer's hard paternalism (to the extent that the employer can be trusted to serve the savers' interests).

4. PRODUCT-USE DISCLOSURE

Regulation of disclosure is another type of soft paternalistic intervention (Bar-Gill [2012] is a useful and comprehensive reference). The literature distinguishes between two subspecies: product-attribute disclosure, which aims to correct biases in how product attributes are perceived and weighted, and product-use disclosure, which aims to help evaluate nonlinear price plans according to an estimated level of consumption (using the consumer's own past behavior or the behavior of other consumers in similar circumstances as an anchor). This section focuses on the latter.

One reason why consumers may need assistance in evaluating price plans is that they tend to make biased predictions of their own future behavior. Bar-Gill (2012) presents a convincing case that underestimating a future taste for immediate gratification may explain why consumers end up paying too much for services such as credit cards or mobile phone services. In particular, he invokes the influential model of DellaVigna and Malmendier (2004), in which firms compete in two-part tariffs for consumers whose intertemporal preferences display so-called hyperbolic discounting. In this section I introduce product-use disclosure into a simple version of this model.

4.1. The Benchmark Model

I begin with the basic DellaVigna-Malmendier model without disclosure. There are three time periods, 0, 1, and 2. At period 0, two firms that provide an identical service (credit, mobile telecommunication, and so on) compete for a measure one of identical consumers. Each firm i simultaneously commits to a nonlinear price scheme t_i: $[0, \infty) \rightarrow \mathbb{R}$, where $t_i(x)$ is the payment the consumer will make to firm i in period 2 conditional on selecting this firm in period 0 and subsequently choosing consumption level x in period 1. Once the consumer chooses a firm in period 0, he is obliged by the firm's price plan in the next two periods. Following DellaVigna and Malmendier (2004), I restrict t_i to take the form of a two-part tariff, namely, $t_i(x) = A_i + p_i x$, and identify it with the pair of parameters (A_i, p_i). Both firms face the same constant marginal cost $c \in (0, 1)$. Firm i's profit is 0 if it is not chosen by the consumer and $t_i(x) - cx$ if the consumer chooses the firm and proceeds to consume x.

Consumers have dynamically inconsistent preferences that take a standard (β, δ) form (also known as hyperbolic discounting) with $\delta = 1$. In period 0, when they face the choice of a price plan, their utility from accepting a two-part tariff (p, A) and proceeding to consume x is $U(x, p, A) = \beta \ln(x + 1) - \beta(px + A)$. In contrast, in period 1, when they are obliged by the price plan (p, A) and choose the consumption quantity x, their utility function changes to $V(x, p, A) = \ln(x + 1) - \beta(px + A)$. The parameter $\beta < 1$ represents the consumers' taste for immediate gratification—namely, it measures how they discount all future utility flows against the current utility flow. From the point of view of period 0, both consumption and payment take place in subsequent periods, and therefore both are discounted by β, but in period 1, consumption is immediate, whereas payment lies in the future, and therefore only the latter is discounted. As a result, the consumers' trade-off between consumption and payment is dynamically inconsistent. To ensure interior solutions in the sequel, I make the simplifying restriction $c < \beta$.

The literature makes a distinction between sophisticated consumers who correctly predict their future preferences and naive consumers who incorrectly believe that their future preferences will be identical to their current preferences. I focus on the case of naive consumers: in period 0, the consumer chooses a price plan that maximizes U under the incorrect assumption that he will choose x in period 1 to maximize U. Firms correctly anticipate that consumers will make consumption decisions in period 1 to maximize V and calculate their profit from any price plan accordingly. The following result, from DellaVigna and Malmendier (2004)—specialized for this example—analyzes Nash equilibrium in the period 0 game between the firms.

Proposition 2. In symmetric Nash equilibrium, each firm offers a two-part tariff t^* given by

$$p^* = \frac{c}{\beta}, \qquad A^* = \frac{(\beta - 1)(1 - c)}{\beta}.$$

This scheme induces an actual consumption quantity $y^* = 1/c - 1$. In period 0, the consumer erroneously predicts that he will consume $y^0 = \beta/c - 1$.

The equilibrium price per unit p^* is above marginal cost. Competitive pressures push the firms' actual equilibrium profits to 0, which means that the lump-sum payment A^* is negative. That is, when accepting t^*, the consumer expects to consume a relatively small quantity y^0, and he is mainly attracted by the lump-sum subsidy. However, having accepted a price plan, his taste for immediate gratification impels him to consume $y^* > y^0$. Since $p^* > c > 0$, while $A^* < 0$, the following inequalities hold:

$$A^* + p^*y^* > A^* + p^*y^0, \qquad \frac{A^* + p^*y^*}{y^*} > \frac{A^* + p^*y^0}{y^0}.$$

That is, both total and per-unit payment end up being higher than the consumer anticipates when accepting the equilibrium price plan. The consumer's total equilibrium payment is $A^* + p^*y^* = 1 - c$. If the consumer correctly predicted his actual period 1 behavior, he would evaluate the equilibrium price plan ex ante at $U(y^*, p^*, A^*) = 1 + \beta(\ln \beta - 1) + \beta(c - \ln c - 1)$.

4.2. Enter the Regulator

Imagine that in response to this situation, a regulator introduces product-use disclosure. For each price plan offered in the market, the regulator mandates the disclosure of the effective price per unit 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 tx^*/x^*. I assume that in period 0 the consumer obediently chooses the firm with the lowest disclosed effective price (with a symmetric tie-breaking rule). Having selected a price plan t given by (p, A), the consumer proceeds in period 1 to choose a consumption quantity x that maximizes $V(x, p, A)$. The interpretation is that the three-period interaction is repeated for many rounds. In each round, there is a new generation of consumers who enter the market and make a once-and-for-all decision, and the disclosure is informed by the historical behavior of previous generations of consumers. This form of disclosure does not provide new information and has no effect on a rational consumer. It is a nudge that manipulates the boundedly rational consumer's method of evaluating price plans.

I wish to emphasize that introducing product-use disclosure into the DellaVigna-Malmendier model is not an arbitrary idea. Indeed, one of the key applications of DellaVigna and Malmendier (2004) is to credit markets, a context in which product-use disclosure is commonly discussed. And as mentioned above, Bar-Gill (2012) presents the DellaVigna-Malmendier model as an explanation for overpayment patterns (of the type captured by proposition 2) in the credit card industry and recommends product-use disclosure as a potential remedy. This section draws logical consequences from this argument. For this purpose, I need to modify the notion of a stable market outcome; I can no longer model the situation as a game, because the market agents' long-run behavior feeds into the firms' payoff function via the disclosed effective price. I therefore rely on the following definition.

Definition 3: Stable Market Outcomes. A triple (x^s, p^s, A^s) is stable if the following conditions hold: (1) $x^s = \text{argmax}_x V(x, p^s, A^s)$, and (2) no firm has an incentive to deviate from (p^s, A^s) to another price plan, given the consumers' rule for choosing a price plan in period 0 (where x^s plays the role of the historical consumption quantity) and their subsequent choice of consumption quantity in period 1 under the price plan they select.

Condition 1 means that in order for the consumption quantity x^s to persist across generations of consumers, it must be optimal for them (according to their preferences in period 1) given the equilibrium price plan (p^s, A^s). Condition 2 reflects the notion that if x^s is a stable consumption level, it becomes the historical quantity that informs the calculation of the effective price. Suppose that a firm deviates from (p^s, A^s) to some other price plan (p, A). Then condition 2 requires either of the following scenarios to be realized: consumers are not attracted to (p, A) because it has a weakly higher disclosed effective price than the prevailing plan (p^s, A^s)—that is, $(A + px^s)/x^s \geq (A^s + p^s x^s)/x^s$—or consumers are attracted to (p, A) because $(A + px^s)/x^s < (A^s + p^s x^s)/x^s$, but the firm's deviation fails to increase its profit, given the way consumers actually choose under (p, A).

Proposition 4. There is a unique stable triple (x^s, p^s, A^s), where $x^s = 1/\beta c - 1$, $p^s = c$, and $A^s = 0$.

Thus, on the face of it, product-use disclosure is effective in the sense that it induces a competitive stable outcome: firms' behavior is reduced to linear, marginal-cost pricing. However, if the motivation behind the regulatory intervention is to reduce overconsumption, then the stable outcome achieves the opposite objective, because $x^s > y^*$. With regard to the consumer's welfare, evaluated according to his ex ante (period 0) perspective, note that $U(x^s, p^s, A^s) = \beta(c - \ln c - 1)$. Given that $\beta < 1$, it is easy to verify that $U(x^s, p^s, A^s) < U(y^*, p^*, A^*)$—that is, consumers' ex ante utility decreases as a result of the regulatory intervention. The consumer's total payment in the stable outcome is $p^s x^s = 1/\beta - c$, which is higher than the total equilibrium payment prior to the intervention.

The intuition for this result is simple. Product-use disclosure causes firms to practice linear pricing, and competitive pressures drive the price per unit down to marginal cost. However, since consumer choice given a price plan is determined by the period 1 utility *V*, firms effectively compete with the "the wrong self"—namely, the consumer's taste for immediate gratification. Fostering this kind of competition is the opposite of what the regulation was meant to do, namely, protect the consumer from the market reaction to his underestimation of this taste, which manifests itself in overconsumption.[7]

Note that my welfare analysis relies on the assumption that consumers have dynamically inconsistent preferences. Suppose that consumers do not have distinct period 0 preferences but simply held the incorrect prior belief that their period 1 preferences would be given by *U* rather than *V*. This alternative assumption would lead to the same positive analysis of the market model (such equivalence would cease to hold if consumers were partially sophisticated—see Spiegler [2011, ch. 5] for a related discussion), but the normative analysis would be different, because there would be no reason to use *U* as a welfare criterion.[8]

5. PRODUCT-ATTRIBUTE DISCLOSURE

Goods and services often have multiple price and quality attributes, and consumers may neglect some of them. For instance, when thinking about the actual cost of a credit card, borrowers may pay more attention to the basic interest rate than to late fees. Headline prices in contracts attract more attention than the small print. Some products have future add-ons that consumers may fail to take into account at the time of purchase (for example, replacement ink cartridges in printers). Can mandated disclosure of product attributes help consumers make better decisions in this context?

An implicit assumption behind attribute disclosure is that once the consumer becomes fully aware of all attributes, he will execute a rational evaluation. However, suppose that the consumer has a fixed attention

7. The stable 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 nonlinear pricing scheme rather than just two-part tariffs (see Spiegler 2011, ch. 2), based on a model by Eliaz and Spiegler (2006). Thus, product-use disclosure ends up simulating an environment in which all restrictions on price plans are lifted.

8. I thank Oren Bar-Gill for pointing out this consideration.

budget, in the sense that he can take into consideration only a subset of attributes. Instead of increasing the consumer's attention budget, disclosure might simply reallocate it among the various attributes. The tendency to neglect relevant attributes can also arise from deeper aversion to 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 an emotionally inconvenient trade-off, and a natural response is to neglect attributes, thus saving the mental cost of dealing with such trade-offs. When one attribute is less salient than another, it is psychologically easier to ignore it. Product-attribute disclosure makes all attributes equally salient, thus placing them on an equal footing in this regard.[9]

What are the implications of product-attribute disclosure under this view of the psychological process that underlies attribute neglect? Will disclosure make consumers better off when equilibrium effects are taken into account? In this section I use the model in Bachi and Spiegler (2014) to capture these considerations. Unlike in previous sections, here none of the results are new; the only novel contribution is their interpretation in terms of product-attribute disclosure. Therefore, I report the results briefly and refer the reader to Bachi and Spiegler (2014) for more general statements of the results and the proofs.

The market (once again) consists of two identical, profit-maximizing firms that compete for a measure one of consumers by playing a simultaneous-move game. Each firm $i = 1, 2$ chooses a product that is fully characterized by a quality vector $(q_i^1, q_i^2) \geq (0, 0)$. (I use the language of quality rather than prices for expositional ease.) Firm i's profit conditional on being chosen is $1 - \frac{1}{2}(q_i^1 + q_i^2)$. I refer to $\overline{q}_i = \frac{1}{2}(q_i^1 + q_i^2)$ as the true quality of the firm's product, where quality is measured in terms of the firm's cost of producing it. 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, both firms would offer quality vectors that maximize u subject to the constraint that true quality is $\overline{q} = 1$ (that is, 0 profits).

Now suppose that quality dimension 2 is shrouded such that all consumers focus entirely on dimension 1 and choose the firm that offers the

9. For psychological evidence of the phenomenon of trade-off avoidance, see Tversky (1972), Payne, Bettman, and Johnson (1993), Luce, Payne, and Bettman (1999), and Anderson (2003).

highest quality along this dimension (with symmetric tie breaking). The following proposition is a special case of the well-known model from Gabaix and Laibson (2006).

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

In equilibrium, competition is effectively restricted to the salient dimension 1, and this enables firms to choose the lowest possible quality along the shrouded dimension 2. Competitive pressure drives quality up along dimension 1, until firms earn 0 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 $\overline{q} = 1$. In fact, the equilibrium strategy maximizes $q^1 - \overline{q}$ subject to the constraint that firms earn nonnegative profits—in this respect, it is maximally misleading.

Imagine a regulator that wishes to curb misleading contracts and responds to this state of affairs by mandating disclosure that will unshroud dimension 2. Furthermore, 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 reallocates their attention budget between the two dimensions, such that every consumer focuses on either of the two attributes with probability $\frac{1}{2}$. As a result, when $q_i^k > q_j^k$ for both $k = 1, 2$, all consumers choose firm i, but when $q_i^1 > q_j^1$ and $q_i^2 < q_j^2$, each firm gets half the consumer population (the case of equality along one or both dimensions is irrelevant for the analysis). What are the equilibrium implications of this intervention?

Proposition 6.[10] When each consumer chooses according to a uniformly drawn single attribute, the game between the two firms has a unique symmetric Nash equilibrium: each firm chooses $q^1 + q^2 = 1$ with probability 1 and draws q^1 uniformly from [0, 1].

Thus, in equilibrium, firms offer products of true average quality $\overline{q} = \frac{1}{2}$, and the breakdown into the two dimensions is random. Each firm earns a profit of $\frac{1}{4}$ in equilibrium. The intuitive reason for firms' ability to earn positive profits is that they can offer a product with high quality in

10. From Bachi and Spiegler (2014).

one dimension and low quality in the other. Some consumers will choose the firm because they happen to focus on the dimension in which the firm is relatively attractive, and yet the firm will be able to make a profit thanks to the low quality it offers on the other dimension.

In equilibrium, no market alternative ever dominates the other. In this sense, consumers always face hard trade-offs in equilibrium. At the same time, equilibrium pricing is less misleading than before the intervention: the difference between perceived and true quality (which is $q^1 - \overline{q}$ and $q^2 - \overline{q}$ with probability $\frac{1}{2}$ each) is uniformly distributed over $[-\frac{1}{2}, \frac{1}{2}]$, compared with the deterministic gap $q^1 - \overline{q} = 1$ in the absence of disclosure. The regulator has succeeded in his mission to curb misleading contracts, but at the cost of lowering the true quality of the products that are offered in equilibrium. This is an example of how two desirable criteria for consumer protection, maximizing quality and preventing misleading contracts, can be mutually conflicting.

This exercise provides another demonstration that having an explicit psychological story behind observed consumer biases matters for the equilibrium analysis of nudging. When consumers seem to be ignoring certain attributes, it matters whether this is a manifestation of simple unawareness or a result of deeper psychological forces such as intrinsic attention deficit or trade-off avoidance. 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, and this lack of predictability weakens competitive pressures.

Behavioral economics models require special caution when it comes to welfare analysis, because of the disconnect between preferences and choices. This is true in the case of the models in Sections 3 and 4, but it holds even more strongly in the present section. The reason is that the consumer's choice procedure—before and after the intervention—is entirely based on ordinal quality rankings; it reveals nothing about how the consumer truly trades off the two dimensions. For this reason, I refrain from making statements about the consumer's true welfare in equilibrium and restrict the discussion to product quality.

6. CONCLUSION

Nudging is appealing because it seems to offer a regulatory free lunch: helping boundedly rational consumers without infringing contractual freedom. This paper demonstrate that, theoretically, equilibrium market responses to nudges can eat away part of this 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 very familiar to economists: when analyzing theoretical consequences of an intervention, it is useful to think about agents' equilibrium reaction in terms of an explicit structural model that accounts for their observed stimulus-response patterns. However, the sense in which the present paper is structural is unusual. Economists normally reserve the term for rational choice models that are explicit about agents' preferences and information. In comparison, the models in this paper are explicit about other mental constructs, such as the ability to make comparisons or predict future tastes.

An important implication of this structural approach is that there can be a big difference between consumers' underlying potential for bias and the amount of bias that is observed. The former is determined by structural features (for example, the parameter λ in Section 2), while the latter is also governed by firms' equilibrium reaction to the regulatory environment. As a result, naive regulatory response to an observed bias may have unintended consequences. I hope that this paper demonstrates that adopting the structural approach enriches the theoretical discussion of consumer protection, even if by its very nature it tends to problematize issues instead of offering easy solutions.

APPENDIX: PROOFS

A1. Proof of 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 x conditional on p.

I begin with a few preliminary observations. First, note that since $\lambda > 0$, firms can secure a strictly positive profit by charging $p > 0$ and playing $x = 1$. Therefore, $p = 0$ is not in support of F. Suppose that F has an atom on any $p < 1$. If $\sigma_p(0) > 0$, then a firm can profitably deviate to a strategy that consists of the price $p' = p - \varepsilon$ and the mixture σ_p over x, where $\varepsilon > 0$ is arbitrarily small. If $\sigma_p(1) = 1$, then a firm can profitably deviate to the pure strategy $(p + \varepsilon, 1)$, where $\varepsilon > 0$

is arbitrarily small. Thus, F is continuous over $p < 1$. Note that I cannot rule out the possibility that F places an atom on $p = 1$ and $\sigma_p(1) = 1$. The reason is that since comparison probability is 0 when $x_1 = x_2 = 1$, deviating to the pure strategy $(1 - \varepsilon, 1)$ is not profitable for an arbitrarily small $\varepsilon > 0$. Finally, the support of F must be an interval $[p_1, 1]$, where $p_1 > 0$—otherwise, if there is a hole (p, p') in the support of F, the strategy consisting of the price p' and the mixture σ_p over x generates a strictly higher payoff than (p, σ_p), which belongs to the support of the equilibrium strategy, a contradiction.

Let $p_H(x)$ and $p_L(x)$ denote the highest and lowest prices p in the closure of the set $\{p \in [p_L, 1] | \sigma_p(x) > 0\}$. Let me show that $p_H(0) \le p_L(0)$. Assume the contrary, namely, that there exist $p^1, p^2 \in [p_L, 1]$ such that $p^2 > p^1$, $\sigma_{p^2}(0) > 0$, and $\sigma_{p^1}(1) > 0$. For each $k = 1, 2$ and $x = 0, 1$, the market share that the strategy (p^k, x) generates, denoted $s(p^k, x)$, is as follows:[11]

$$s(p^k,\ x) = \frac{\lambda}{2}\left\{1 - \int_{p_L}^{p^k}\left[\sum_y \sigma_p(y)\left(1 - \frac{1}{2}x - \frac{1}{2}y\right)\right]dF(p)\right\}$$
$$+\left(1 - \frac{\lambda}{2}\right)\int_{p^k}^{1}\left[\sum_y \sigma_p(y)\left(1 - \frac{1}{2}x - \frac{1}{2}y\right)\right]dF(p).$$

It is now straightforward to verify that it is impossible that $s(p^1, 1) \ge s(p^1, 0)$ and $s(p^2, 0) \ge s(p^2, 1)$, because F assigns positive probability to the interval (p^1, p^2).

Suppose that $p_L(1) < 1$. I have just shown that $\sigma_p(1) = 1$ for every $p \in [p_L(1), 1]$. Since comparison probability is 0 when $x_1 = x_2 = 1$, if a firm deviates from a price $p \in [p_L(1), 1]$ to the pure strategy $(1, 1)$, its market share does not change, and hence its payoff increases. Therefore, $p_L(1) = 1$, such that $\sigma_p(0) = 1$ for almost every $p \in [p_L, 1]$. If F does not place an atom on $p = 1$, this means that $x = 0$ is played with probability 1, and in this case a firm that charges a price close to 1 will strictly prefer to deviate to $x = 1$. Thus, it must be the case that F places an atom on $p = 1$. Let μ denote the size of this atom. The following equality must hold:

$$s(1,\ 1) = \frac{\lambda}{2} \times \left[\mu + \frac{1}{2}(1 - \mu)\right] = \frac{\lambda}{2} \times \mu + \left(1 - \frac{\lambda}{2}\right) \times \frac{1}{2}\mu = \lim_{\varepsilon \to 0} s(1 - \varepsilon,\ 0).$$

Otherwise, there would be a profitable deviation either from (1, 1) to (1, 0) or from $(p, 0)$ to $(p, 1)$ for some p sufficiently close to 1. Thus, $\mu = \lambda/2$ such that a firm's payoff from (1, 0) is $(\lambda/2)(\frac{1}{2} + \lambda/4)$. Since this is the equilibrium payoff, it is the payoff from $(p, 1)$ for every $p \in [p_L, 1)$. Thus, I can write

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

11. To simplify the notation, I ignore the possibility that F has an atom on pk—to incorporate an atom I would have to replace F with left or right limits of F—without changing the argument.

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

A2. Proof of Proposition 2

As DellaVigna and Malmendier (2004) show, in this environment the symmetric equilibrium two-part tariff (p^*, A^*) maximizes consumers' perceived ex ante net utility $U(x^u, p, A)$ subject to the zero-profit condition

$$A + (p - c)x^v = 0,$$

where $x^u = \text{argmax}_x U(x, p, A)$ and $x^v = \text{argmax}_x V(x, p, A)$. Since V is concave and twice differentiable in x, x^v is simply given by the first-order condition $V'(x, p, A) = 1/(x + 1) - \beta p = 0$; hence, $x^v = 1/\beta p - 1$ (as long as $\beta p \leq 1$). Likewise, x^u is given by the first-order condition $U'(x, p, A) = \beta/(x + 1) - \beta p = 0$; hence, $x^u = (1/p) - 1$ (as long as $p \leq 1$). To find p^*, I thus need to find the value of p that maximizes

$$\ln(x^u + 1) - px^u + (p - c)x^v = \ln\frac{1}{p} - p\left(\frac{1}{p} - 1\right) + (p - c)\left(\frac{1}{\beta p} - 1\right).$$

Solving this maximization problem gives $p^* = c/\beta < 1$. Inserting this value into the above equations for x^u, x^v, and A completes the characterization of equilibrium.

A3. Proof of Proposition 4

Let me first show that the triple $(x^s, p^s, A^s) = (1/\beta c - 1, c, 0)$ is stable. Suppose that a firm deviates to some other price plan (p, A). For consumers to choose (p, A) over (p^s, A^s), it must be the case that

$$A + p\left(\frac{1}{\beta c} - 1\right) < c\left(\frac{1}{\beta c} - 1\right).$$

Consumers who select (p, A) will subsequently choose x to maximize $V(x, p, A)$; hence, $x = 1/\beta p - 1$. For the deviation to be profitable for the firm, it must be that

$$A + (p - c)\left(\frac{1}{\beta p - 1}\right) > 0.$$

The two inequalities are clearly contradictory.

Let me now show that there is no other stable pair (x, p, A). Once again, stability requires $x = 1/\beta p - 1$. I now show that stability requires firms to earn 0 profits; that is, $A + (p - c)(1/\beta p - 1) = 0$. If $A + (p - c)(1/\beta p - 1) < 0$, a firm can deviate to (p, A'), $A' > A$. Consumers will not choose the firm, and so it will make 0 profits; hence, the deviation is profitable. If $A + (p - c)(1/\beta p - 1) > 0$, a firm can deviate to (p, A), where $\varepsilon > 0$ is arbitrarily small. All consumers will select the firm because it obviously has a lower effective price, and they will proceed to choose $x = 1/\beta p - 1$ because the price per unit has not changed; hence, the increase in market share outweighs the slight loss in the profit per customer.

Suppose that the triple $(x^s, p^s, A^s) = (1/\beta c, c, 0)$ is unstable. Then I can find (z, p, A) such that

$$A + p\left(\frac{1}{\beta c - 1}\right) < c\left(\frac{1}{\beta c - 1}\right), \qquad A + (p - c)z > 0,$$

where $z = (1/\beta p - 1)$. Then the following inequality must hold:

$$(p - c)\left(1 - \frac{1}{\beta p}\right) < (p - c)\left(1 - \frac{1}{\beta c}\right),$$

and it can be easily verified that no p can satisfy it. Therefore, (x^s, p^s, A^s) is stable.

REFERENCES

Anderson, Christopher. 2003. The Psychology of Doing Nothing: Forms of Decision Avoidance Result from Reason and Emotion. *Psychological Bulletin* 129:139–67.

Armstrong, Mark. 2008. Interactions between Competition and Consumer Policy. *Competition Policy International* 4:97–147.

———. 2015. Search and Ripoff Externalities. *Review of Industrial Organization* 47:273–302.

Armstrong, Mark, and John Vickers. 2012. Consumer Protection and Contingent Charges. *Journal of Economic Literature* 50:477–93.

Bachi, Benjamin, and Ran Spiegler. 2014. Buridanic Competition. Unpublished manuscript. Tel Aviv University, School of Economics, Tel Aviv.

Baker, Tom, and Peter Siegelman. 2014. Behavioral Economics and Insurance Law: The Importance of Equilibrium Analysis. Research Paper No. 14-1. University of Pennsylvania, Institute for Law and Economics, Philadelphia.

Bar-Gill, Oren. 2012. *Seduction by Contract: Law, Economics, and Psychology in Consumer Markets*. Oxford: Oxford University Press.

Barr, Michael S., Sendhil Mullainathan, and Eldar Shafir. 2008. Behaviorally Informed Financial Services Regulation. White paper. New America Foundation, Washington, DC.

Beshears, John, James Choi, David Laibson, and Brigitte Madrian. 2012. Simplification and Saving. *Journal of Economic Behavior and Organization* 95:130–45.

Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin. 2003. Regulation for Conservatives: Behavioral Economics and the Case for "Asymmetric Paternalism." *University of Pennsylvania Law Review* 151:1211–54.

Carlin, Bruce I. 2009. Strategic Price Complexity in Retail Financial Markets. *Journal of Financial Economics* 91:278–87.

Chioveanu, Ioana, and Jidong Zhou. 2013. Price Competition with Consumer Confusion. *Management Science* 59:2450–69.

de Clippel, Geoffroy, Kfir Eliaz, and Kareen Rozen. 2013. Competing for Consumer Inattention. Discussion Paper No. 9553. Centre for Economic Policy Research, London.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. Contract Design and Self-Control: Theory and Evidence. *Quarterly Journal of Economics* 119:353–402.

Duarte, Fabian, and Justine Hastings. 2012. Fettered Consumers and Sophisticated Firms: Evidence from Mexico's Privatized Social Security Market. Working Paper No. 18582. National Bureau of Economic Research, Cambridge, MA.

Eliaz, Kfir, and Ran Spiegler. 2006. Contracting with Diversely Naive Agents. *Review of Economic Studies* 73:689–714.

Ellison, Glenn. 2006. Bounded Rationality in Industrial Organization. Pp. 2:142–74 in *Advances in Economics and Econometrics: Theory and Applications*, edited by Richard Blundell, Whitney K. Newey, and Torsten Persson. Ninth World Congress. New York: Cambridge University Press.

Gabaix, Xavier, and David Laibson. 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics* 121:505–40.

Grubb, Michael D. 2015. Consumer Inattention and Bill-Shock Regulation. *Review of Economic Studies* 82:219–57.

Grubb, Michael D., and Matthew Osborne. 2015. Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock. *American Economic Review* 105:234–71.

Handel, Benjamin. 2013. Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *American Economic Review* 103:2643–82.

Heidhues, Paul, and Botond Koszegi. 2010. Exploiting Naivete about Self-Control in the Credit Market. *American Economic Review* 100:2279–2303.

Huck, Steffen, and Jidong Zhou. 2011. Consumer Behavioural Biases in Competition: A Survey. Report commissioned by the Office of Fair Trade, London.

Iyengar, Sheena, Gur Huberman, and Wei Jiang. 2004. How Much Choice Is Too Much? Determinants of Individual Contributions in 401K Retirement Plans. Pp. 83–95 in *Pension Design and Structure: New Lessons from Behavioral Finance*, edited by Olivia S. Mitchell and Stephen P. Utkus. Oxford: Oxford University Press.

Kamenica, Emir, Sendhil Mullainathan, and Richard Thaler. 2011. Helping Consumers Know Themselves. *American Economic Review* 101:417–22.

Luce, Mary, John Payne, and James Bettman. 1999. Emotional Trade-off Difficulty and Choice. *Journal of Marketing Research* 36:143–59.

Madrian, Brigitte C., and Dennis F. Shea. 2001. The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *Quarterly Journal of Economics*

116:1149–87.

Mullainathan, Sendhil, Joshua Schwartzstein, and William Congdon. 2012. A Reduced-Form Approach to Behavioral Public Finance. *Annual Review of Economics* 4:511–40.

Muris, Timothy. 2002. The Interface of Competition and Consumer Protection. Paper presented at Fordham Corporate Law Institute's 29th annual International Antitrust Law and Policy conference, New York.

Payne, John, James Bettman, and Eric Johnson. 1993. *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.

Piccione, Michele, and Ran Spiegler. 2012. Price Competition under Limited Comparability. *Quarterly Journal of Economics* 127:97–135.

Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.

Spiegler, Ran. 2006a. Competition over Agents with Boundedly Rational Expectations. *Theoretical Economics* 1:207–31.

———. 2006b. The Market for Quacks. *Review of Economic Studies* 73:1113–31.

———. 2011. *Bounded Rationality and Industrial Organization*. Oxford: Oxford University Press.

Sunstein, Cass. 2014. *Why Nudge? The Politics of Libertarian Paternalism*. New Haven, CT: Yale University Press.

Thaler, Richard, and Cass Sunstein. 2003. Libertarian Paternalism. *American Economic Review* 93:175–79.

———. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.

Tversky, Amos. 1972. Elimination by Aspects: A Theory of Choice. *Psychological Review* 79:281–99.

Varian, Hal R. 1980. A Model of Sales. *American Economic Review* 70:651–59.