American Economic Journal: Microeconomics 2016, 8(1): 1–23
http://dx.doi.org/10.1257/mic.20140138

Exploitative Innovation†

By Paul Heidhues, Botond Kőszegi, and Takeshi Murooka*

We analyze innovation incentives when firms can invest either in increasing the product's value (value-increasing innovation) or in increasing the hidden prices they collect from naive consumers (exploitative innovation). We show that if firms cannot return all profits from hidden prices by lowering transparent prices, innovation incentives are often stronger for exploitative than for value-increasing innovations, and are strong even for non-appropriable innovations. These results help explain why firms in the financial industry (e.g., credit-card issuers) have been willing to make innovations others could easily copy, and why these innovations often seem to have included exploitative features. (*JEL* D21, G21, L11, L25, O31)

A growing theoretical literature in behavioral economics investigates how firms use hidden fees—e.g., overdraft fees for bank accounts and late fees and high interest payments for credit cards—to exploit naive consumers. This research raises a fundamental question: where do the hidden fees come from? Inventing a new way to exploit naive consumers, much like inventing any novel product feature, presumably requires innovation, and existing research has not investigated the incentives for such "exploitative innovation." Indeed, since many exploitative features—especially in financial products—seem to be in easily copyable contract terms, from a classical perspective the incentives to invent them are unclear.

In this paper, we analyze the incentives for exploitative innovation in a market for potentially deceptive products, and contrast them with the often-studied incentives for making product improvements that consumers value. Section I introduces our model, which consists of a simultaneous-move price-competition stage modeled after Gabaix and Laibson (2006) and Heidhues, Kőszegi, and Murooka (2014), and a preceding innovation stage. At the price-competition stage, firms selling perfect substitutes each set a transparent upfront price as well as an additional price,

*Heidhues: European School of Management and Technology, Schlossplatz 1, 10178 Berlin, Germany (e-mail: Paul.Heidhues@esmt.org); Kőszegi: Central European University, Nádor utca 9. 1051 Budapest, Hungary (e-mail: botondkoszegi@gmail.com); Murooka: University of Munich, Ludwigstr. 28 Rgb, 80539 Munich, Germany (e-mail: takeshi.murooka@econ.lmu.de). The setup of this paper was developed in our earlier working paper "The Market for Deceptive Products," which we split into two papers. The companion paper is titled "Inferior Products and Profitable Deception." We are grateful to Özlem Bedre-Defolie, Stefano DellaVigna, Drew Fudenberg, Michael Grajek, Michael Grubb, Péter Kondor, Nicolas Melissas, Klaus Schmidt, Antoinette Schoar, and seminar and conference audiences for insightful comments and feedback. Heidhues and Kőszegi also gratefully acknowledge financial support from the European Research Council (Starting Grant #313341). This paper was revised while Heidhues visited the Institute of Advanced Studies at the Central European University (CEU), which he thanks for its hospitality.

†Go to http://dx.doi.org/10.1257/mic.20140138 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

and unless at least one firm decides to costlessly unshroud (i.e., educate consumers about) additional prices, naive consumers ignore these prices when making purchase decisions. To capture the notion that in some markets, such as credit cards and banking services, firms cannot return all profits from later charges by lowering initial charges, we deviate from most existing work and posit that there is a floor on the upfront price. (In Section IB, we provide one possible microfoundation that endogenously gives rise to a floor on the upfront price.) We assume that whenever a deceptive equilibrium—wherein all firms shroud (i.e., do not educate consumers about) additional prices—exists at the pricing stage, firms play that continuation equilibrium.[1]

To investigate incentives at the innovation stage in the simplest possible manner, we assume that only one firm, firm 1, can make innovations. Firm 1 can invest either in exploitative innovation—increasing the maximum additional price—or in value-increasing innovation—increasing the product's value—and other firms observe its innovation decision. We consider both appropriable innovations (i.e., innovations other firms cannot copy) and non-appropriable innovations (i.e., innovations other firms can fully copy), as well as in-between cases.

In Section II, we characterize innovation incentives when the price floor is not binding—a condition we argue may hold for some commonly invoked examples of deceptive products, including hotel rooms and printers. In this case, the incentive for both an exploitative and a value-increasing innovation is based on the "appropriable part" of the innovation. If firm 1 increases its product value by \$10 relative to others, it can charge almost \$10 more than competitors and still capture the entire market, earning a profit of almost \$10 per consumer. And if firm 1 figures out a way to charge a \$10 higher additional price than others, it can charge slightly lower prices than competitors, capture the entire market, and make \$10 more ex post, again earning a profit of almost \$10 per consumer. This can help explain why firms have developed some appropriable exploitative practices, such as a proprietary technology that prevents printer users from buying non-brand printer cartridges, in industries with a non-binding price floor. But because firms have no incentive to make non-appropriable innovations, and the primary tools for exploitation often seem to be easily copyable contract terms, the extent of deception in these industries may in the end be limited.

In Section III, we characterize innovation incentives when the price floor is binding—a condition we argue holds for consumer financial products such as credit cards and banking services. Now, the incentive is often stronger for exploitative than for value-increasing innovation. To understand the most important part of the logic, suppose for a moment that independently of the innovation no firm wants to unshroud additional prices, or (in a simplified variant of our model) unshrouding is impossible. Since the floor on the upfront price is binding, firm 1 always benefits from an exploitative innovation by being able to increase its margin. In contrast,

[1] Whenever such an equilibrium exists, it is the most plausible one for two reasons. Most importantly, it is then the unique equilibrium in the variant of our model in which unshrouding is costly, no matter how small the cost is. In addition, in most of our settings a deceptive equilibrium Pareto dominates an unshrouded-prices (and hence zero-profit) equilibrium from the perspective of the firms, and in some settings it does so strictly.

firm 1 cannot benefit from a non-appropriable value-increasing innovation, so its incentive to make such an innovation is zero. The threat of unshrouding—the threat that a competitor, unable to compete on the upfront price, prefers to unshroud and compete on the additional price, eliminating profitable deception—only makes the contrast between these incentives greater. Because acquiring an exploitative innovation increases competitors' profits from shrouding and thereby lowers their motive to unshroud, the innovation may enable profitable deception in the industry. This increases firm 1's incentive to make exploitative innovations, and means that firm 1 may be willing or eager to share its innovation with competitors. But because a non-appropriable value-increasing innovation increases competitors' motive to unshroud by raising the profits from selling the product transparently, the possibility of unshrouding lowers a firm's incentive to make such an innovation.

To complete our analysis, we explore the incentives for fully or partially appropriable value-increasing innovations, showing that even these are stronger for products that should not survive in the market in the first place. An appropriable value-increasing innovation steals the consumers of other firms, allowing firm 1 to capture the entire market. In a socially valuable industry where consumers value the product above production costs however, this induces competitors to unshroud, dampening firm 1's incentive to innovate. In contrast, in a socially wasteful industry where consumers value the product below production cost, a firm that unshrouds cannot go on to profitably sell its product, so no firm has an incentive to unshroud. As a result, firm 1's incentive to innovate is stronger.

Our model implies substantial incentives for exploitative innovation. From a welfare perspective, all of this spending is of course a pure social waste. To make matters worse, exploitative innovation can enable profitable deception in the industry to the detriment of consumers, and—because it is only with sufficiently high additional prices that consumers can be fooled into buying a socially wasteful product—may facilitate the emergence of a socially wasteful industry. The situation is especially bleak in an industry with a binding price floor, where innovation incentives are tilted in favor of exploitative innovations over value-increasing innovations, and where a firm may have an incentive to make even non-appropriable exploitative innovations. These insights can help explain why firms in many financial industries have been willing to make contract innovations with deceptive features that others could easily copy, and raise the general concern that resources are directed disproportionately toward these kinds of innovations.

Our paper builds on a theoretical literature in behavioral economics that explores how firms use hidden fees or otherwise take advantage of mistakes in consumer decisions.[2] To our knowledge, however, neither this literature nor the extensive classical literature on research and development has investigated firms' incentives to make exploitative innovations. In fact, our message that a firm may be eager to invest in socially wasteful non-appropriable exploitative innovation contrasts with a

[2] See, for instance, DellaVigna and Malmendier (2004), Eliaz and Spiegler (2006), Gabaix and Laibson (2006), Spiegler (2006a, 2006b), Grubb (2009), Heidhues and Kőszegi (2010), Piccione and Spiegler (2012), Grubb (2015), Ko (2012), and Spiegler (2011) for a textbook treatment.

prevalent theme in the classical literature that firms selling substitutes often underinvest in non-appropriable innovations (e.g., Reinganum 1989).

I. Basic Model

A. *Setup*

We introduce our model of innovation in a market for potentially deceptive products. The game has two stages, an innovation stage and a price-setting stage. The price-setting stage is a variant of the model in Heidhues, Kőszegi, and Murooka (2014), and we begin by describing this stage. There are $N \geq 3$ firms competing for a unit mass of naive consumers who value firm n's product at $v_n > 0$ and are looking to buy at most one item. Firms simultaneously set upfront prices f_n and additional prices a_n, and decide whether to costlessly unshroud the additional prices. The highest possible additional price firm n can charge is $\bar{a}_n > 0$, and a consumer buying product n has to pay both prices f_n and a_n. If all firms shroud, consumers make purchase decisions believing that the total price of product n is f_n. If at least one firm unshrouds, all consumers become aware of all additional prices, and hence make purchase decisions based on the total prices $f_n + a_n$. Crucially, we deviate from much of the literature and impose that firms face a floor on the upfront price: $f_n \geq \underline{f}$.[3] We assume that $\underline{f} \leq v_n$ for all n, so that in a shrouded market consumers are willing to buy from a firm with an upfront price at the floor. We do not impose a floor on the total price.

Each firm's cost of providing the product is $c > 0$. We assume that $v_n + \bar{a}_n > c$ for all n; a firm with $v_n + \bar{a}_n < c$ cannot profitably sell its product, so without loss of generality we can think of it as not participating in the market. We make two simple tie-breaking assumptions. First, consumers go to a highest-quality firm when indifferent.[4] Second, if all firms shroud and a subset of firms with the same quality choose an upfront price at the floor, these firms split their demand in proportion to shares $s_n \in [0, 1)$.

We now turn to describing the innovation stage. To identify innovation incentives in a transparent manner, we assume that only one firm, firm 1, can make innovation investments. Prior to the innovation, each firm has a product with value v and maximum additional price $\bar{a}$. Then, firm 1 chooses whether or not to invest in innovation, with all firms observing its decision. We consider separately two types of innovation. An "exploitative innovation" costs I_a and increases the maximum additional price firm 1 can charge by Δa and the maximum additional price firm $n \neq 1$ can charge by $\Delta a'$, where $0 \leq \Delta a' \leq \Delta a$. This formulation allows us to consider the two extreme cases often studied in the literature, appropriable innovations (which competitors cannot copy: $\Delta a' = 0$) and non-appropriable innovations (which competitors can fully copy: $\Delta a' = \Delta a$), as well as in-between cases. At the

[3] Armstrong and Vickers (2012), Grubb (2015), and Ko (2012) also analyze models with variants of our price-floor assumption.

[4] This allows a firm to price a lower-quality competitor out of the market by offering the same deal—something it could do anyway by offering a minimally better deal—ensuring the existence of a pure-strategy equilibrium. This simplifies some of our proofs, but does not affect the logic of our results.

price-setting stage, $\bar{a}_1 = \bar{a} + \Delta a$ and $\bar{a}_n = \bar{a} + \Delta a'$ for $n \neq 1$ if the innovation takes place; otherwise $\bar{a}_n = \bar{a}$ for all n. Analogously, a value-increasing innovation costs I_v and increases consumers' valuation of product 1 by Δv and their valuation of product $n \neq 1$ by $\Delta v'$, where $0 \leq \Delta v' \leq \Delta v$.[5]

We look for subgame-perfect Nash equilibria in the game played between firms, imposing—as is standard in the industrial-organization literature—that no firm charges a total price below its marginal cost. In addition, we assume that if a deceptive continuation equilibrium—wherein all firms shroud additional prices—exists in the pricing subgame, firms play that continuation equilibrium. Because no firm has an incentive to shroud if at least one firm unshrouds, there is always an unshrouded-prices continuation equilibrium, which results in Bertrand price competition. When a deceptive continuation equilibrium exists, however, it is more plausible than the unshrouded-prices continuation equilibrium for two main reasons. Most importantly, in any situation in which we use this assumption, the deceptive equilibrium is the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is.[6] In addition, in most of our settings a deceptive equilibrium Pareto dominates an unshrouded-prices equilibrium from the perspective of the firms, and in some settings it does so strictly. Using these refinements, we characterize investment incentives by identifying the maximum investment costs I_a^* and I_v^* below which firm 1 is willing to make the investment of each type.

To conclude the setup of our model, we briefly mention how the additional price and exploitative innovation map to economic settings that have been invoked as conducive to deception. For credit cards, additional prices can result because consumers underestimate the amount of interest or fees they will pay. An issuer can invent marketing strategies to induce consumers to borrow more, or introduce new types of fees that consumers do not fully understand. For banking services, additional prices can arise because consumers underappreciate fees for overdraft protection or other services. A bank can invent new hidden fees that consumers will likely run into. For printers, additional prices are generated by consumers' failure to anticipate high cartridge prices. A design change that makes generic cartridges less compatible with the firm's printer increases the maximum cartridge price the firm can charge.

In each of the above applications, firms have some scope to educate consumers about additional prices. For instance, a firm can make cartridge prices or credit-card fees more salient. Indeed, Stango and Zinman (forthcoming) document that directing consumers' attention to overdraft fees reduces overdrafts for up to two years. Nevertheless, our assumption on unshrouding—that a single firm can educate all

[5] While we interpret value-increasing innovations as increasing the product's true value to consumers, the same results hold for innovations, advertisement, and other investments that merely increase the perceived value—with the investment's social value of course being lower in this case than for true value-increasing innovations. For example, a mutual-fund prospectus outlining an investment philosophy may fool consumers into believing that there is a dependable way to beat the market, increasing the perceived value of the fund.

[6] The proof of Proposition 5 in Heidhues, Kőszegi, and Murooka (2012b) establishes this claim for any situation in which consumers value firms' products equally and the price floor is binding, a condition that applies to Proposition 2. For Part (ii) of Proposition 1 and for Proposition 3, we establish the claim formally in the proof of the corresponding proposition. In Proposition 4, we use our equilibrium refinement only for the subgame following no innovation, and in this case consumers value products equally and the price floor is binding. In Part I of Proposition 1, we do not use the refinement.

consumers about all additional prices at no cost—is in the context of most real settings unrealistically extreme, especially if incurring the additional prices depends partly on the consumer's own behavior. When presenting our results, we discuss what happens when unshrouding is impossible.

B. *Microfoundation for the Floor on the Upfront Price*

This subsection provides a foundation for the floor on the upfront price based on Heidhues, Kőszegi, and Murooka (2012a). This microfoundation is a discrete, simplified variant of Ellison's (2005) insight (developed in the model of add-on pricing) that firms may be reluctant to cut upfront prices because these cuts disproportionately attract less profitable consumers. Other potential microfoundations for the floor are extensively discussed in Heidhues, Kőszegi, and Murooka (2012b).

Suppose that in addition to the naive consumers described above, there is a proportion $1 - \alpha$ of "arbitrageurs" in the market. Arbitrageurs respond to money-making opportunities, and can avoid paying the additional price by incurring cost e. For simplicity, we assume that arbitrageurs enter the market only when this is strictly profitable, and can buy from multiple firms. Lemma 1 shows that if the fraction of arbitrageurs is sufficiently large, a price floor arises.

LEMMA 1: *Suppose that $\alpha \bar{a}_n < e + c$ for all n. Then, the sets of Nash equilibria when firms face no restriction on $f_n (f_n \in \mathbb{R})$ and when firms face the restriction $f_n \geq -e$ are the same. Furthermore, in any Nash equilibrium arbitrageurs do not buy.*

The intuition is simple. If a firm sets an upfront price below $-e$ when others set a higher price, it attracts not only all naive consumers, but also all unprofitable arbitrageurs to itself. To avoid this outcome, firms act as if they were facing a price floor of $-e$.

Since many or most products are easy to get and dispose of, a price floor of zero or somewhat below zero is likely to apply to most products. Hence, we will assume throughout this paper that $\underline{f} = 0$. In some cases, e could be negative—and the price floor therefore positive—if the product has an alternative use that arbitrageurs value (e.g., the scanner feature of a multifunction printer).[7]

Of course, the way we model the threat of arbitrage—that a firm does not attract any arbitrageurs if it sets an upfront price above a bright-line number, but faces a flood of arbitrageurs or sophisticated consumers if it sets a lower upfront price—is extreme. The intuitions for our main qualitative results on profitable deception require only that firms are less willing to cut the upfront price when shrouding than the transparent total price when unshrouding, so that they make higher profits with shrouded than with unshrouded additional prices. Even without the stark arbitrageur behavior we impose above, this is the case if—similar to Ellison (2005)—firms disproportionately attract less profitable customers when cutting their upfront price, but not when cutting their transparent total price.

[7] In this case, assuming that arbitrageurs benefit from getting additional products is unrealistic. Slightly modifying our proof, however, establishes that Lemma 1 applies even if arbitrageurs can get only one unit.

II. Non-Binding Price Floor

Largely as a benchmark for our main results in the next section, we first consider innovation incentives when the floor on the upfront price is not binding. More precisely, we suppose that $\underline{f} \leq c - (\bar{a} + \Delta a)$—i.e., that a firm cannot make positive profits when setting an upfront price equal to the floor. This condition may hold for some commonly invoked examples of deceptive products, such as printers and hotel rooms (e.g., Hall 1997, Gabaix and Laibson 2006). For instance, the marginal cost of a hotel room (c) is likely to be nontrivial, and the additional amount a hotel can extract from the minibar, telephone, and other add-on services $(\bar{a} + \Delta a)$ is limited. Hence, with a price floor of around \$0, the above inequality is satisfied.

Proposition 1 identifies innovation incentives in this case:

PROPOSITION 1: *Suppose* $\underline{f} \leq c - (\bar{a} + \Delta a)$. *Then,*

(*i*) (*Value-Increasing Innovation*). $I_v^* = \Delta v - \Delta v'$.

(*ii*) (*Exploitative Innovation*). $I_a^* = \Delta a - \Delta a'$.

Part (i) of Proposition 1 says that firm 1's incentive to make a value-increasing innovation is based on the "appropriable part" of the innovation—the extent to which the innovation increases the value of its product above that of competitors' products. As a simple example, suppose firms' cost is \$100, innovation increases the value of firm 1's product from \$200 to \$220 and the value of other firms' products from \$200 to \$210, and the maximum additional price—which firms actually charge in any deceptive continuation equilibrium—is \$50. Similarly to the logic of Lal and Matutes (1994), classical switching-cost models, and many existing behavioral-economics models with naive consumers, firms compete aggressively for ex post-profitable consumers, and bid down the upfront price to \$100 − \$50 = \$50. Absent the innovation, therefore, firm 1 cannot sell its product above an upfront price of \$50, so it earns zero net profits. If it innovates, however, firm 1 can charge an upfront price slightly below \$60 and attract all consumers, generating total revenue of nearly \$110 with the additional price included. And because other firms are charging a total price equal to marginal cost, no firm can profitably unshroud and attract consumers. Hence, firm 1 earns a profit of \$10 per consumer.

Part (ii) of Proposition 1 says that similarly to its incentive to make a value-increasing innovation, firm 1's incentive to make an exploitative innovation is equal to the appropriable part of the innovation. In this case, however, firm 1's competitive advantage derives not from offering a better product to consumers, but from better exploiting consumers. Continuing with the above example, suppose the innovation increases firm 1's maximum additional price to \$70 and other firms' maximum additional price to \$60. Because of the higher additional price they can charge, other firms are now willing to bid down the upfront price to \$40. Even so, firm 1 can offer a slightly lower upfront price, attract all consumers, and again earn a revenue of nearly \$110 with its higher additional price included. In other words, the profitability of exploitative and value-increasing innovations is exactly the same: a value-increasing innovation allows a firm to raise its total price above competitors' and still keep consumers; and an exploitative innovation allows a firm to lower its price below competitors' and still make profits.

An example consistent with the prediction that a firm will make appropriable exploitative innovations is the printer industry. Hall (1997) describes a number of strategies, including questionable "artistic" cartridge design patents, printer-head patents, and perpetual design modifications, that generate no consumer value but help printer manufacturers control the cartridge market and thereby cash in on naive consumers. Nevertheless, with a non-binding price floor a firm's incentive to make exploitative innovations is no greater than its incentive to make value-increasing innovations, and in particular it has no incentive to make non-appropriable exploitative innovations. Because the primary tools for deception often seem to be contract terms that tie the consumer to the firm and induce her to pay supra-normal fees ex post, and such contract innovations are typically easy to copy, the extent of deception in industries with a non-binding price floor may in the end be limited. While there are other plausible explanations, the contrast between this observation and those in the next section is consistent with a difference in industry practices between hotels and credit cards. Hotels could charge high fees for a variety of contingencies, such as cancellations or modifications, and possibly do so in ways that consumers do not fully anticipate. Yet unlike credit-card issuers, hotels do not seem very focused on expanding such sources of revenue.

III. Binding Price Floor

A. *Main Results*

We now turn to the main goal of our paper: analyzing innovation incentives when the price floor is binding. More precisely, we assume throughout this section that $\underline{f} > c - \bar{a}$; a firm can (if it is able to charge a sufficiently high additional price) make positive profits when setting an upfront price equal to the floor. This case describes a number of consumer financial products, including credit cards, bank accounts, and actively-managed mutual funds. For instance, the marginal cost of setting up a credit-card account for a consumer (c) is quite low, while credit card companies make substantial amounts in hidden fees ($\bar{a}$ is large), so with $\underline{f} = 0$ the above inequality is easily satisfied.[8]

To analyze innovation incentives, we first identify conditions under which a deceptive equilibrium exists in the pricing subgame when firms' products are equally valuable. Recall that we assume a deceptive equilibrium will be played in any such situation.

[8] For the US credit card market, Evans and Schmalensee (2005) estimate that the average cost of opening a new account, including all marketing and processing cost, is about $72. And as argued, for instance, by Ausubel (1991), credit card companies make large ex post profits on charges consumers do not anticipate. Similarly, based on data from the United Kingdom, Armstrong and Vickers (2012) emphasize that banks make substantial amounts in overdraft fees.

LEMMA 2 (Equilibrium in the Pricing Subgame): *Suppose* $\underline{f} > c - \bar{a}_n$ *and* $v_n = v'$ *for all n. If*

$$s_n(\underline{f} + \bar{a}_n - c) \geq v' - c \qquad \text{(SC)}$$

holds for all n, then a deceptive continuation equilibrium exists. In any deceptive continuation equilibrium, $f_n = \underline{f}$ *and* $a_n = \bar{a}_n$ *for all n, and firms earn positive profits. If* (SC) *is violated for some n, then in any continuation equilibrium prices are unshrouded with probability one, consumers buy at a total price of c, and firms earn zero profits.*

The intuition for why firms might earn positive profits despite facing Bertrand-type price competition is in two parts. First, firms make positive profits from the additional price, and to obtain these ex post profits each firm wants to compete for consumers by offering better upfront terms. But the price floor prevents firms from competing away all profits from the additional price by lowering the upfront price. Second, since a firm cannot compete for consumers by cutting its upfront price, there is pressure for it to compete on the additional price—but because competition in the additional price requires unshrouding, it is an imperfect substitute for competition in the upfront price. A firm that unshrouds and undercuts competitors tells consumers not only that its product is the cheapest, but also that the product is more expensive than they thought. This surprise may lead consumers not to buy, in which case the unshrouding firm can attract consumers only if it cuts the total price by a substantial margin. Since this may not be worth it, the firm may prefer not to unshroud.

In Heidhues, Kőszegi, and Murooka (2014), we show that profitable deception is likely to be more pervasive in socially wasteful industries where the value consumers derive from the product is below marginal cost than in socially valuable industries where the value is strictly above marginal cost. If the product is socially wasteful, a firm that unshrouds cannot profitably sell its product, so with no firm ever wanting to unshroud a profitable deceptive equilibrium always exists. But if the product is socially valuable, a firm that would make sufficiently low profits from deception can earn higher profits from unshrouding and capturing the entire market; so if there is such a firm, only a non-deceptive, zero-profit equilibrium exists.

We now turn to innovation incentives. Proposition 2 states our results for non-appropriable innovations, showing that they are stronger for exploitative than for value-increasing innovations:

PROPOSITION 2 (Non-Appropriable Innovations):

(*i*) (*Exploitative.*) *Suppose* $\underline{f} > c - \bar{a}$, $\Delta a = \Delta a'$, *and all firms satisfy* (*SC*) *for* $\bar{a}_n = \bar{a} + \Delta a, v' = v$. *If all firms satisfy* (*SC*) *for* $\bar{a}_n = \bar{a}, v' = v$, *then* $I_a^* = s_1 \Delta a$. *If some firm n does not satisfy* (*SC*) *for* $\bar{a}_n = \bar{a}$, $v' = v$, *then* $I_a^* = s_1(\underline{f} + \bar{a} + \Delta a - c) > s_1 \Delta a$.

(*ii*) (*Value-Increasing.*) *Suppose* $\underline{f} > c - \bar{a}$, $\Delta v = \Delta v'$, *and all firms satisfy* (*SC*) *for* $\bar{a}_n = \bar{a}, v' = v$. *If all firms satisfy* (*SC*) *for* $\bar{a}_n = \bar{a}, v' = v + \Delta v$, *then* $I_v^* = 0$. *If at least one firm n violates* (SC) *for* $\bar{a}_n = \bar{a}, v' = v + \Delta v$, *then* $I_v^* < 0$.

To start understanding the logic behind Proposition 2, suppose that with or without the innovation a deceptive continuation equilibrium obtains (i.e., (SC) holds). Clearly, in this case innovation incentives are the same as in a setting where unshrouding is impossible in the first place. The first statements in the two parts of the proposition say that in this case, firm 1 is willing to spend resources on non-appropriable exploitative innovation, but not on non-appropriable value-increasing innovation. Since any exploitative innovation (be it appropriable or non-appropriable) does not lead to a decrease in the upfront price, it simply increases firm 1's markup by Δa, and hence increases firm 1's profits by exactly $s_1 \Delta a$. But because a non-appropriable value-increasing innovation can increase neither one's market share nor one's markup, firm 1 has no incentive to invest in it.

The key difference between exploitative and value-increasing innovations that generates the different incentives for them is that the former acts on the firm's markup, while the latter acts on consumers' willingness to pay (WTP). With a binding price floor, a firm always benefits from a markup-side innovation: even if the innovation is non-appropriable, it does not lead to increased competition on the price, and hence increases the profits of the firm. But in our model with inelastic demand, a firm cannot benefit from a non-appropriable WTP-side innovation, since it does not increase the firm's market share. It is worth noting that a cost-reducing innovation is a markup-side innovation in the above sense, so that a firm would have an incentive to make a non-appropriable cost-reducing innovation as well.

We now consider the implications of the possibility that innovation affects firms' willingness to go along with a deceptive equilibrium, as captured in (SC). As Part (i) of Proposition 2 states, if firms cannot maintain a deceptive equilibrium without the exploitative innovation, the innovation enables profitable deception in the industry, so firm 1's willingness to pay for the innovation is equal to its full post-innovation profits—a potentially huge incentive to innovate. Intuitively, competitors who are not very good at imposing additional prices gain little from deception and hence may want to deviate from it, threatening the deceptive equilibrium and thereby firm 1's profits. To eliminate such a threat, firm 1 would like to teach competitors how to better exploit consumers. And as Part (ii) shows, it may be the case that firms can maintain a deceptive equilibrium without but not with a value-increasing innovation, so that firm 1's willingness to pay for the innovation is negative. Intuitively, an increase in v does not affect profits when firms shroud, but—by increasing the amount consumers are willing to pay for a transparent product—does increase the profits a firm can gain from unshrouding. As a result, firm 1 may be willing to spend money to avoid an increase in v. Hence, the threat of unshrouding increases the incentive to make non-appropriable exploitative innovations—and, equivalently, increases a firm's willingness to share exploitative innovations—and decreases the incentive to make non-appropriable value-increasing innovations.

The message of Proposition 2 that firms might be willing to make investments in non-appropriable innovations, and that such innovations are likely to be exploitative rather than socially valuable, seems consistent with how some consumer financial products have developed recently. Many features that have been identified by researchers as deceptive—such as teaser rates, high fees for certain patterns of product use, and difficult-to-understand payment schedules involving large future payments—are contract innovations that seem easy to copy, and that in fact have been copied quickly by competitors. For credit cards, this is especially the case since many issuers participate in information exchange that allows them to easily observe each other's practices.[9]

To complete our analysis, we consider fully or partially appropriable value-increasing innovations. We distinguish between socially wasteful and socially valuable industries, beginning with the former one.

PROPOSITION 3 (Value-Increasing Innovation in Socially Wasteful Industries): *Suppose* $\underline{f} > c - \bar{a}, v + \Delta v < c,$ *and* $\Delta v > \Delta v'$. *Then,* $I_v^* = \left[(1 - s_1)(\underline{f} + \bar{a} - c)\right] + \left[\Delta v - \Delta v'\right]$.

Proposition 3 implies that firm 1's willingness to pay for fully or partially appropriable value-increasing innovations in a socially wasteful industry—that is, for products that should not be in the market in the first place—is quite high: it is greater than in the corresponding classical setting (where it would clearly be zero); it is greater than the increase in the relative value of firm 1's product $(\Delta v - \Delta v')$; and (because I_v^* is bounded away from zero), it is nontrivial even for vanishingly small product improvements. Firm 1's willingness to pay, I_v^*, derives from two sources. First, as captured in the first term, the innovation attracts the consumers of all competitors to firm 1, and firm 1 benefits from this even at pre-innovation market prices. Second, as captured in the second term, because the innovation improves firm 1's product more than competitors' products, firm 1 can increase the upfront price without losing consumers, further increasing its profits. Although firm 1 makes these extra profits by pricing competitors out of the market, with the industry being socially wasteful competitors do not unshroud in response.

We next consider socially valuable industries.

PROPOSITION 4 (Value-Increasing Innovation in Socially Valuable Industries): *Suppose* $\underline{f} > c - \bar{a}$, $v > c$, $\Delta v > \Delta v'$, *and* (SC) *holds for all* n *when* $\bar{a}_n = \bar{a}$, $v' = v$. *Then,* $I_v^* = \left[\Delta v - \Delta v'\right] - s_1(\underline{f} + \bar{a} - c)$.

To help interpret Proposition 4, suppose for a moment that unshrouding is impossible. Then, firm 1's willingness to pay for a partially appropriable value-increasing innovation is—by the same logic and formal argument—the same as in Proposition 3. Proposition 4 implies that with the threat of unshrouding, firm 1's willingness to pay for the same type of innovation is lower. In a socially valuable industry, any partially appropriable innovation must lead to unshrouding; otherwise,

[9] In particular, Argus is an information-exchange service that collects individual-level account data from credit card issuers and, based on this data, relays information on current practices to other issuers. The information Argus collects includes fee assessment practices, strategies for balance generation, financial performance, and payment behavior. Argus emphasizes that it has detailed information on "virtually every US consumer credit card." See http://www.argusinformation.com/eng/our-services/syndicated-studies/credit-card-payment-study/us-credit-card-payments-study/.

firm 1 would be able to price competitors out of the market while setting a high total price, and competitors would respond by unshrouding and profitably undercutting this total price. Hence, innovation leads firm 1 to lose its positive profits from deception. This loss dampens firm 1's incentive to innovate.

B. *Extension: Downward-Sloping Demand*

Our model above assumes inelastic demand for the product—that increasing the product's value does not increase total demand—thereby abstracting from one benefit of value-increasing innovation. We now consider downward-sloping demand. Let $D(f)$ be the demand curve induced by the distribution of naive consumers' valuations prior to any innovation. We suppose that $D(f)$ is continuously differentiable, log-concave, $D(\underline{f}) > 0$, and that there is a choking price (i.e., an f' such that $D(f) = 0$ for $f \geq \overline{f}'$). We compare non-appropriable exploitative innovations with non-appropriable value-increasing innovations that increase the valuation of each consumer by Δv, inducing a post-innovation demand curve $D(f - \Delta v)$. For simplicity, we assume that unshrouding is impossible; we have confirmed that much like in our basic model, the possibility of unshrouding increases the incentive to make exploitative innovations and decreases the incentive to make value-increasing innovations. Let f^m be firm 1's monopoly upfront price when it charges $a_1 = \bar{a}$, which is unique since demand is log-concave. Then, one has:

PROPOSITION 5 (Non-appropriable Innovation with Downward-Sloping Demand): *If $c - \bar{a} < \underline{f} < f^m$ and $\Delta a = \Delta v$ with both sufficiently small, then $I_a^* > I_v^*$.*

The proposition states that if $\underline{f}$ is lower than the monopoly price—an arguably weak assumption—then firm 1's incentive to make a small non-appropriable exploitative innovation is strictly greater than its incentive to make a similarly small non-appropriable value-increasing innovation. An exploitative innovation that raises $\bar{a}$ by one unit increases the firm's markup, enabling it to gain profits on every infra-marginal unit it sells. Increasing consumers' valuations by one unit has the same effect on demand as lowering the upfront price by one unit, and hence raises profits by inducing extra sales on the margin. Below the monopoly price, the infra-marginal effect of a price increase is greater than the marginal loss in sales, so the former effect outweighs the latter.

While Proposition 5 is stated for small innovations, we can identify its broader implications by thinking of large innovations as a sequence of small innovations. If the floor is below the monopoly price at any intermediate step in any such sequence, then $I_a^* > I_v^*$ holds even for larger $\Delta a = \Delta v$.

IV. Conclusion

Given our emphasis that exploitation requires innovation, it would be interesting to investigate the dynamics of how exploitation appears and spreads in an industry. One possible scenario suggested by our theory is the following. With the industry initially in a non-deceptive situation, one firm invents and starts offering a product with shroudable features (e.g., it adds costly overdraft protection to its bank accounts). Because neither consumers nor competitors were aware of this product, it starts off being shrouded. At this point, competitors must decide whether to unshroud the product or to adopt it in their product lines. Our theory suggests that competitors' preference is to adopt the deceptive product.

An obvious question raised by our paper's negative picture of innovation incentives is whether and how a policymaker can improve market outcomes. This question gives rise to a number of difficult issues that require substantial further research. For instance, reliably identifying whether an innovation is value-increasing or exploitative, or whether a product is socially wasteful, seems extremely difficult, and it is important to develop general methods with which regulators can make these determinations. Furthermore, even if it was possible to determine whether a contract or product innovation is value-increasing, a policy that requires explicit approval would likely reduce innovation. Finding regulatory responses that lower the incentives for exploitative but not for value-increasing innovation is an important open policy question. But even before such questions have been satisfactorily answered, our results caution against the common intuition that the reduction of innovation incentives is necessarily disadvantageous.

Appendix

PROOF OF LEMMA 1:

We establish that $f_n < -e$ is strictly dominated for firm n, implying that the sets of Nash equilibria when f_n is unrestricted ($f_n \in \mathbb{R}$) and when f_n is restricted ($f_n \geq -e$) are the same.

We first note that setting $f_n + a_n < 0$ is strictly dominated for firm n, because then all arbitrageurs take up firm n's product and firm n earns negative profits. Hence, we suppose from now on that $f_n + a_n \geq 0$.

Given $f_n + a_n \geq 0$, arbitrageurs take up firm n's product (and pay the cost e to avoid a_n) if and only if $f_n < -e$ and $a_n > e$. We show that such pricing is again strictly dominated for firm n. With these prices, since all arbitrageurs take up firm n's product and avoid a_n, firm n's total profit is at most

$$(1-\alpha)(f_n - c) + \alpha(f_n - c + \bar{a}_n) < -e - c + \alpha\bar{a}_n < 0.$$

Therefore, $f_n < -e$ and $a_n > e$ are strictly dominated by $f_n > v, a_n = 0$.

Finally, notice that arbitrageurs do not buy if $f_n \geq -e$ and $f_n + a_n \geq 0$. Hence, in any Nash equilibrium, arbitrageurs do not enter the market. ■

PROOF OF PROPOSITION 1:

Part (i). It is easy to check that the following is an equilibrium in the pricing subgame: all firms shroud with probability one and set an additional price of $\bar{a}$, firm 1 sets an upfront price of $c - \bar{a} + \Delta v - \Delta v'$, all other firms set an upfront price of $c - \bar{a}$, and firm 1 gets the entire market and earns a profit of $\Delta v - \Delta v'$.

We next argue that firm 1 earns at least $\Delta v - \Delta v'$ in any equilibrium of the pricing subgame. Recall that by assumption, no firm prices below marginal cost. Thus, if firm 1 unshrouds and charges a total price of $c + \Delta v - \Delta v'$, all consumers weakly prefer firm 1's product and our tie-breaking assumption implies that firm 1 serves the entire market, earning $\Delta v - \Delta v'$.

We will now argue that firm 1 earns no more than its relative advantage $\Delta v - \Delta v'$. Suppose otherwise. Since firm 1 earns more than its relative advantage, it must do so (in expectation) for all but a set of measure zero of total prices it charges; hence there exists an $\epsilon > 0$ such that firm 1 charges a total price above $c + \Delta v - \Delta v' + \epsilon$ for some $\epsilon > 0$ with probability 1. Any firm $k \neq 1$ must earn positive profits bounded away from zero; otherwise it could deviate, unshroud and offer a total price of $c + \epsilon/2$, thereby offering a better deal to consumers than firm 1 and hence win with positive probability and earn positive profits. Furthermore, since the equilibrium outcome does not coincide with that of the corresponding standard Bertrand game, all firms must shroud with positive probability.

Let $\hat{t}_k$ be the supremum of the total price distribution firm $k \neq 1$ charges; and let $\hat{t}_1$ be that of firm 1. Define the quality-adjusted maximum of these suprema as $\hat{t} = \max\{\hat{t}_1 - \Delta v, \hat{t}_k - \Delta v'\}$. Note that firm k cannot charge this quality-adjusted total price with positive probability when unshrouding; if it did, it would lose to firm 1 with probability one, contradicting the fact that it must earn positive profits with any price it charges with positive probability. Furthermore, firm k cannot charge this quality-adjusted price with positive probability when shrouding. If it did, it must have positive market share and it can do so only if all other firms shroud. But then firm k must set $a_k = \bar{a}$ to maximize profits, and hence it offers a contract $(\hat{t} + \Delta v' - \bar{a}, \bar{a})$ with positive probability. For this contract to have positive market share, firm 1 must shroud and set its upfront prices strictly above $\hat{t} + \Delta v - \bar{a}$ with positive probability, and by the definition of $\hat{t}$ firm 1 at the same time must set the additional price strictly below $\bar{a}$. But this is not a best response: because $\hat{t}_1 - \Delta v \leq \hat{t}$, firm 1 could always keep the total price distribution fixed and charge the maximal additional price $\bar{a}$ so that its upfront prices always lie weakly below $\hat{t} + \Delta v - \bar{a}$, and firm 1 could thereby strictly increase its market share holding the total price distribution fixed. We conclude that firm k cannot charge a quality-adjusted price of $\hat{t}$ with positive probability. But then, firm 1 also cannot do so when unshrouding as it would have zero market share. Furthermore, firm 1 cannot charge $\hat{t}$ with positive probability when shrouding: if it did, firm k would have to charge its upfront prices weakly above $\hat{t} + \Delta v' - \bar{a}$ with positive probability when shrouding for firm 1 to have positive market share. Also, firm k's additional prices associated with these upfront prices must be strictly below $\bar{a}$ because we showed that firm k cannot charge a quality-adjusted price of $\hat{t}$ with positive probability, i.e., $\hat{t}_k - \Delta v' < \hat{t}$. But then, firm k could always keep its total price distribution fixed and charge the maximal additional price $\bar{a}$, thereby ensuring that its upfront prices are strictly below $\hat{t} + \Delta v' - \bar{a}$, increasing its market share while holding the total price fixed. We conclude that no firm charges the highest quality-adjusted price with positive probability.

We now show that as $\epsilon \to 0$, the market shares of both firm 1 and firm k when charging a quality-adjusted price in the interval $(\hat{t} - \epsilon, \hat{t})$ go to zero. This will imply that their profits go to zero, contradicting the fact that they must earn positive profits (bounded away from zero) for all but a set of measure zero of prices.

The above statement is immediate if the firm in question unshrouds; hence, for a sufficiently small $\epsilon > 0$ it must be that firms 1 and k almost always shroud when they set total quality-adjusted prices in the interval $(\hat{t} - \epsilon, \hat{t})$. Suppose that firm 1 shrouds on this interval but its market share does not approach zero. With positive probability, firm k must shroud and set upfront prices weakly above $\hat{t} + \Delta v' - \bar{a}$, while setting total prices strictly below $\hat{t} + \Delta v'$. But then, firm k could always keep its total price distribution fixed and charge the maximal additional price $\bar{a}$, increasing its market share for a set of total prices with positive probability. The argument for why the market share of firm k must go to zero when shrouding and charging quality-adjusted total prices in $\left(\hat{t} - \epsilon, \hat{t}\right)$ is analogous. We conclude that firm 1 earns its relative advantage in every equilibrium of the pricing subgame, which leads to Part (i) of the proposition.

Part (ii). It is easy to check that the following constitutes an equilibrium in the pricing subgame: all firms shroud and set $(f_n, a_n) = (c - \bar{a}_n, \bar{a}_n)$, firm 1 gets the entire market, and earns a profit of $\Delta a - \Delta a'$. We next argue that profits are the same in any continuation equilibrium in which firms shroud with probability one and no firm charges below marginal costs. Note that any equilibrium outcome of the pricing subgame in which all firms shroud with probability one must also be an equilibrium outcome of the corresponding pricing game in which shrouding is impossible. Consider this pricing game. Clearly, all firms that sell with positive probability to consumers must set the maximum additional price, and the additional price of a firm that does not sell to consumers is inconsequential. Hence, we can think of the pricing game as one in which all firms set the maximal additional price. This game is equivalent to a Bertrand-competition game in which firm 1 has cost $c - (\bar{a} + \Delta a)$ and all other firms have cost $c - (\bar{a} + \Delta a')$. It is well-known that in any equilibrium of this game in which no firm charges below marginal cost, firm 1 earns a profit equal to $\Delta a - \Delta a'$. Therefore, in any equilibrium of the pricing subgame in which shrouding occurs with probability one, firm 1 earns a profit equal to $\Delta a - \Delta a'$.

We next show that if firms face an unshrouding cost $\eta > 0$, then in any continuation equilibrium shrouding occurs with probability 1. The proof is by contradiction and has three steps.

Step 1: *No firm unshrouds the additional price with probability one:* If a firm unshrouds with probability one, all consumers become sophisticated and hence buy from the firm with the lowest total price $f + a$. Hence by a standard Bertrand competition argument, firms make zero gross profits (not counting the unshrouding cost) following unshrouding. Then, the firm that chooses to unshroud makes negative net profits (counting the unshrouding cost)—a contradiction.

Step 2: *All firms earn positive profits:* Suppose first that at least two firms unshroud with positive probability. Any firm that does so earns positive gross profits after unshrouding. Then, any firm can shroud but mimic a pricing strategy of a firm with unshrouding—i.e., choose the same distribution over total prices conditional on unshrouding—and can earn positive profits if the other firm unshrouds. Now suppose only one firm unshrouds with positive probability. Still, by the previous argument, all other firms earn positive expected profits, so we are left to show that this firm does. If it is firm 1 and $\Delta a > \Delta a'$, this is clear: since other firms shroud with probability 1, firm 1 can guarantee itself positive profits by setting $f_1 = c - \bar{a} - \Delta a' - \epsilon, a_1 = \bar{a} + \Delta a$ for sufficiently small $\epsilon > 0$. To prove the other cases, suppose the firm mixing between shrouding and unshrouding is firm j. Since $N \geq 3$, there is another firm that shrouds with probability one, earns positive expected profits, and has the same cost and additional price as firm j. Then, firm j can earn positive expected profits by imitating the strategy of this firm.

Step 3: *Main step:* Let $\hat{t}$ be the supremum of the total prices set by any firm conditional on unshrouding. This supremum exists because unshrouding is costly and consumers would not buy at prices exceeding v when unshrouding occurs. Note that it cannot happen that two firms set $\hat{t}$ with positive probability: if this was the case, since a firm earns positive gross profits at that price, it would have an incentive to undercut the other firm. Now suppose that firm j achieves the supremum. At total prices sufficiently close to $\hat{t}$ set by firm j, its market share when unshrouding is not sufficient to cover the unshrouding cost, so that it must earn positive profits when all other firms shroud. Let $(\hat{f}, \hat{a})$ be the associated upfront and additional prices at one such total price for which firm j earns its expected equilibrium profits and plays a best response (which it must do for almost all pricing pairs).

We first show that $j \neq 1$. Suppose, toward a contradiction, that $j = 1$. Consider a deviation by firm 1 in which it shrouds and sets $\hat{f}' = \hat{f} - (\bar{a} + \Delta a - \hat{a}) - \frac{\eta}{2}$, $\hat{a}' = \bar{a} + \Delta a$. This weakly increases firm 1's profit if another firm unshrouds. Furthermore, if all other firms shroud, then with this deviation firm 1's demand is at least as high if it shrouds as if it unshrouds and sets $(\hat{f}, \hat{a})$. Hence, since unshrouding is costly, the deviation increases profits, a contradiction. Thus, we conclude that $j \neq 1$.

Now, let $\tilde{f}$ be the supremum of the upfront prices set by firms other than j conditional on shrouding. Then, for firm j to earn its equilibrium expected profits with $(\hat{f}, \hat{a})$, it must be that $\tilde{f} \geq \hat{f}$. We first rule out equality. Suppose, toward a contradiction, that $\tilde{f} = \hat{f}$. If a firm other than j sets $\tilde{f}$ with positive probability, firm j prefers to undercut; if not, firm j earns zero profits, in either case a contradiction. Hence, $\tilde{f} > \hat{f}$.

Let $(\tilde{f}'', \tilde{a}'')$ be some prices set by firms other than j when shrouding that satisfy $\tilde{f}'' > \hat{f}$, and suppose it is firm k that sets these prices. Let $\pi_k > 0$ be firm k's equilibrium expected profits. For $\tilde{f}''$ sufficiently close to $\tilde{f}$, firm k cannot earn π_k when shrouding occurs. Hence, firm k must earn positive expected profits when unshrouding occurs. Hence, $\tilde{f}'' + \tilde{a}'' \leq \hat{t}$.

Suppose first that $\tilde{f}'' + \tilde{a}'' = \hat{t}$. Then, firm k can earn positive profits when unshrouding occurs only if $\hat{t}$ is set with positive probability. But then firm k prefers to undercut, a contradiction.

Hence, we are left to consider the case $\tilde{f}'' + \tilde{a}'' < \hat{t}$. Consider a deviation by firm k in which it sets $(\hat{f} - \epsilon, \tilde{a}'' + \tilde{f}'' - \hat{f} + \epsilon)$. Since $j \neq 1$, this is feasible so long as $\tilde{f}'' + \tilde{a}'' + \epsilon < \hat{f} + \hat{a}$, which holds for $\hat{f} + \hat{a}$ sufficiently close to $\hat{t}$ and ϵ that is sufficiently small and positive. This deviation does not affect profits if unshrouding occurs, and increases profits if shrouding occurs, a contradiction.

We therefore conclude that according to our equilibrium selection criterion *shrouding occurs with probability one.* Hence, firm 1 earns a profit equal to $\Delta a - \Delta a'$ in the pricing subgame, immediately implying Part (ii) of the proposition. ∎

PROOF OF LEMMA 2:

Note first that if some firm unshrouds with probability one, all other firms are indifferent between shrouding and unshrouding. Thus, an equilibrium in which unshrouding occurs with probability one always exists. In any such unshrouded-prices equilibrium, consumers observe and take the additional price into account, so that our game reduces to a standard Bertrand game in which consumers' willingness to pay is v'. Hence, in any unshrouded-prices equilibrium, consumers buy the product only if $v' \geq c$. In case $v' > c$, standard Bertrand-competition arguments imply that all consumers buy the product at total price of c.

Now consider deceptive equilibria, i.e., equilibria in which shrouding occurs with probability one. In case firm n has a positive probability of sales in equilibrium, it must set $a_n = \bar{a}_n$ as otherwise it could increase its profits conditional on a sale by increasing a_n without affecting the probability of selling. We show that if (SC) holds for all n, then there is a deceptive equilibrium in which all firms set $(\underline{f}, \bar{a}_n)$. With these prices, consumers are indifferent between firms, so firm n gets market share s_n and therefore earns a profit of $s_n(\underline{f} + \bar{a}_n - c)$. For this to be an equilibrium, no firm should want to unshroud additional prices and undercut competitors. Once a firm unshrouds, consumers will be willing to pay exactly v' for its product, so that firm n cannot make profits exceeding $v' - c$ by unshrouding and capturing the entire market. Hence, unshrouding is unprofitable for firm n if (SC) holds.

Next, we prove that firms set $\underline{f}$ with probability one in any deceptive equilibrium. The proof is akin to a standard Bertrand-competition argument. Take as given that all firms shroud with probability 1, and that firm n sets the additional price $\bar{a}_n$. Note that by setting $f_n = \underline{f}$, firm n can guarantee itself a profit of $s_n(\underline{f} + \bar{a}_n - c_n) > 0$. As a result, no firm will set $f_n > v'$, because then no consumer would buy from it. Take the supremum $\overline{f}$ of the union of the supports of firms' upfront price distributions. We consider two cases. First, suppose that some firm sets $\overline{f}$ with positive probability. In this case, all firms have to set $\overline{f}$ with positive probability; otherwise, a firm setting $\overline{f}$ would have zero market share and hence zero profits with probability one. Then, we must have $\overline{f} = \underline{f}$; otherwise, a firm could profitably deviate by moving the probability mass to a slightly lower price. Second, suppose that no firm sets $\overline{f}$ with positive probability. Let firm n's price distribution achieve the supremum $\overline{f}$. Then, as f_n approaches $\overline{f}$, firm n's expected market share, and hence expected profit approaches zero—a contradiction.

To complete the proof, we establish by contradiction that if (SC) is violated for some firm, then in any equilibrium additional prices are unshrouded with probability one. The proof proceeds in three steps.

Step 1: *All firms earn positive profits.* If shrouding occurs with positive probability, then firms must earn positive profits: if all competitors shroud the additional prices, firm n can guarantee itself positive profits by shrouding and offering $\left(\underline{f}, \bar{a}_n\right)$, which attracts consumers since $v' \geq \underline{f}$ and makes positive profits since $\underline{f} + \bar{a}_n > c_n$.

Step 2: *All firms choose the upfront price $\underline{f}$ whenever they shroud.* Consider the supremum of the total price $\hat{t}_n$ set by firm n when unshrouding, and let $\hat{t} = \max_n\{\hat{t}_n\}$. Note that there exists at most one firm that sets this price with positive probability; if two did, then either could increase profits by moving this probability mass to slightly below $\hat{t}$. Let n be the firm that puts positive probability mass on $\hat{t}$ if such a firm exists; otherwise, let n be a firm that achieves this supremum. For firm n to be able to earn its equilibrium profits for prices at or close to $\hat{t}$, all competitors of n must set a total price weakly higher than $\hat{t}$ with positive probability. By the definition of $\hat{t}$, this means that all competitors of n charge a total price weakly higher than $\hat{t}$ with positive probability when shrouding.

First, suppose all firms other than n set a total price strictly higher than $\hat{t}$ with positive probability. Because each firm $n' \neq n$ makes zero profits when unshrouding occurs, it must make positive profits when shrouding occurs. In addition, since it only makes profits when shrouding occurs, it sets the additional price $\bar{a}_{n'}$ with probability one. Take the supremum of firms' upfront prices $\overline{f}'$ conditional on the total price being strictly higher than $\hat{t}$. Because consumers do not buy the product if the upfront price is greater than v' and firms must earn positive profits by Step 1, $\overline{f}' \leq v'$. Note that $\overline{f}' + \bar{a}_{n'} > \hat{t}$ for any $n' \neq n$.

We now show that $\overline{f}' = \underline{f}$ by contradiction. Suppose $\overline{f}' > \underline{f}$. If two or more firms set $\overline{f}'$ with positive probability when shrouding, each of them wants to minimally undercut—a contradiction.

If only one firm n' sets $\overline{f}'$ with positive probability, then firm n' has zero market share both when unshrouding occurs and when shrouding occurs and some firm other than n' sets a total price strictly greater than $\hat{t}$. Because firm n' earns positive profits by Step 1 and is the only firm that sets $\overline{f}'$ with positive probability conditional on the total price being strictly higher than $\hat{t}$, every firm except for n' sets its upfront price strictly higher than $\overline{f}'$ and its total price weakly lower than $\hat{t}$ when shrouding with positive probability. Suppose first $n' = n$. Then, there exists a firm $n'' \neq n$ that shrouds and sets an upfront fee $f_{n''} > \overline{f}'$, $a_{n''} \leq \hat{t} - f_{n''}$ with positive probability. Since $\overline{f}' + \bar{a}_m > \hat{t}$ for any $m \neq n$, $\bar{a}_{n''} > \hat{t} - \overline{f}'$. Then, firm n'' can increase its profits by decreasing all prices $f_{n''} > \overline{f}'$ to $\overline{f}'$ and by increasing its additional price, holding the total price constant—a contradiction. Next, suppose $n' \neq n$. Then, firm n shrouds, sets $f_n > \overline{f}'$ with positive probability and charges an additional price $a_n \leq \hat{t} - f_n$ with probability one when charging these upfront prices. For almost all of these upfront prices, firm n must earn strictly positive profits when shrouding occurs; otherwise firm n could unshroud with probability one and guarantee positive profits when all rivals shroud and charge a total price above $\hat{t}$. Thus, firm n' shrouds and sets $f_n' \geq f_n > \overline{f}'$, $a_{n'} \leq \hat{t} - f_{n'}$ with positive probability. Since $\overline{f}' + \bar{a}_m > \hat{t}$ for any $m \neq n$, firm n' can increase its profits by decreasing all prices $f_{n'} > \overline{f}'$ to $\overline{f}'$ and increasing its additional price, holding the total price constant—a contradiction.

If no firm sets $\overline{f}'$ with positive probability, there exists firm n' that for any $\epsilon > 0$ sets upfront prices in the interval $\left(\overline{f}' - \epsilon, \overline{f}'\right)$ with positive probability. As $\epsilon \to 0$, the probability of firm n' charging the highest upfront price conditional on shrouding and the total price being strictly higher than $\hat{t}$ goes to one. Therefore, the profits go to zero with probability one when unshrouding occurs or when shrouding occurs and some other firm sets a total price strictly greater than $\hat{t}$. Now follow the same steps as in the previous paragraph to derive a contradiction. Thus, we establish that $\overline{f}' = \underline{f}$.

Because $\overline{f}' = \underline{f}$, each firm $n' \neq n$ sets an upfront price of $\underline{f}$ with probability one conditional on its total price being strictly higher than $\hat{t}$. Hence, $\underline{f} + \bar{a}_{n'} \geq \hat{t}$ for any $n' \neq n$. We now show that whenever shrouding, any firm $n' \neq n$ does not set upfront prices strictly above $\underline{f}$ with positive probability. Suppose by contradiction that firm n' sets prices above $\underline{f}$ with positive probability when shrouding. As n' sets $\underline{f}$ with probability one when charging a total price strictly above $\hat{t}$, the associated additional price must almost always satisfy $a_{n'} \leq \hat{t} - f_{n'}$ when shrouding and setting the upfront price strictly above $\underline{f}$. Since n' sets upfront prices strictly above $\underline{f}$ with positive probability when shrouding, there exists an upfront price $g' > \underline{f}$ such that firm n' sets prices above g' with positive probability. There cannot be a competitor whose upfront price when shrouding falls on the interval $\left[\underline{f}, g'\right]$ with positive probability; otherwise, firm n' could increase its profits by decreasing all prices above g' to $\underline{f}$ and by increasing its additional price holding the total price constant. But then, firm n' can raise its upfront price from $\underline{f}$ to g' and increase profits—a contradiction. Thus, any firm $n' \neq n$ sets the upfront price $\underline{f}$ with probability one when shrouding.

Now suppose that firm n charges an upfront price strictly above $\underline{f}$ when shrouding with positive probability. Then it can only earn profits when unshrouding occurs and hence must almost always charge a total price less than or equal to $\hat{t}$ when shrouding. But if it unshrouds and sets the same prices, it would also earn profits when all rivals shroud and set a price above $\hat{t}$, thereby strictly increasing its profits—a contradiction. Hence, firm n also must set $\underline{f}$ with probability one when shrouding.

Second, suppose some firm $n' \neq n$ sets its total price equal to $\hat{t}$ with positive probability. Then, by the above argument no other firms set total price $\hat{t}$ with positive probability. Take the supremum of firms' upfront prices $\overline{f}'$ conditional on the total price being greater than or equal to $\hat{t}$. The remainder of the proof is the same as above.

Step 3: *Additional prices are unshrouded with probability one.* Suppose not. Then, each firm chooses to shroud with positive probability. Take the infimum of total prices $\underline{t}$ set by any firm when shrouding. We consider two cases. First, suppose $\underline{t} \leq v'$. Take a firm that achieves the infimum. By Step 1, this firm earns positive profits. For any $\epsilon > 0$, take total prices below $\underline{t} + \epsilon$ of the firm. By unshrouding and setting $\underline{t} - \epsilon$, the firm decreases its profits by at most 2ϵ when one or more other firms unshroud, but discretely increases its market share if all other firms shroud. Hence, for sufficiently small $\epsilon > 0$ this is a profitable deviation—a contradiction. Second, suppose $\underline{t} > v'$. Take firm n that violates (SC). By Step 2, firm n charges the upfront price $\underline{f}$ whenever it shrouds. Note that firm n's profits are zero when a rival unshrouds, and its profits are at most $s_n(\underline{f} + \bar{a}_n - c_n)$ when shrouding occurs. But then, deviating and setting a total price equal to v' is profitable because conditional on others shrouding firm n would earn $v' - c_n > s_n(\underline{f} + \bar{a}_n - c_n)$. ■

PROOF OF PROPOSITION 2:

We first prove Part (i). In the subgame following an innovation by firm 1, the shrouding condition holds for all firms by assumption, and thus firm 1 earns $s_1(\underline{f} + \bar{a} + \Delta a - c)$ in this case. In the subgame in which firm 1 did not innovate, Lemma 2 implies that firm 1 earns $s_1(\underline{f} + \bar{a} - c)$ if the shrouding condition holds for all firms and zero otherwise. In the former case the innovation increases firm 1's profits by $s_1 \Delta a$, in the latter case by $s_1(\underline{f} + \bar{a} + \Delta a - c)$, which is strictly greater than $s_1 \Delta a$ because $\underline{f} + \bar{a} > c$.

We now prove Part (ii). Lemma 2 implies that firm 1 earns zero profits in the pricing subgame whenever some firm violates the shrouding condition. If all firms satisfy the shrouding condition, firm 1 earns $s_1(\underline{f} + \bar{a} - c)$, which is positive and independent of v. The result, hence, follows from the fact that an increase in v either does not affect whether the shrouding condition holds or leads to a violation of the shrouding condition for some firm. ■

PROOF OF PROPOSITION 3:

Because we consider a socially wasteful industry, all firms must earn zero profits if unshrouding occurs. Therefore, if we introduce positive unshrouding costs $\eta > 0$, shrouding occurs with probability one in equilibrium. Our equilibrium selection criterion, thus, implies that we can focus on equilibria in which shrouding occurs with probability one in the pricing subgames (and we establish the existence of such equilibria below).

We now solve for the equilibria of the subgames following firm 1's innovation decision. Absent innovation, the shrouding condition is satisfied as we are in a socially wasteful industry. Hence, a deceptive equilibrium exists and (using our selection criterion) firm 1 therefore earns $s_1(\underline{f} + \bar{a} - c)$.

Next consider the subgame following a decision to innovate by firm 1. We first establish that there exists an equilibrium in which firm 1 offers the contract $(\underline{f} + \Delta v - \Delta v', \bar{a})$ with probability one, and all firms $n \neq 1$ offer the contract $(\underline{f}, \bar{a})$; in this equilibrium all consumers are indifferent between firm 1 and its best competitor and following our tie-breaking rule buy firm 1's product. Since unshrouding yields zero profits in a socially wasteful industry, it is immediate that there exist no deviation for any firm $n \neq 1$ that yields positive profits. If firm 1 deviates and unshrouds or shrouds and sets a higher upfront price, it earns zero profits. And since firm 1 captures entire market share, it cannot benefit from lowering its upfront or additional price. Hence, firm 1 also plays a best response.

To complete the proof, we show that in any pricing subgame following innovation in which firms shroud with probability one, firm 1 sets $\left(\underline{f} + \Delta v - \Delta v', \overline{a}\right)$ with probability one and gets the entire market, so that our equilibrium-selection criterion selects such an equilibrium. This means that firm 1 earns $\underline{f} + \Delta v - \Delta v' + \overline{a} - c$ if it innovates and $s_1\left(\underline{f} + \overline{a} - c\right)$ if it does not. I_v^* is the difference between these two profit levels.

To prove the above, we begin by showing that in any equilibrium of the pricing subgame following innovation firms $n \neq 1$ earn zero profits. Suppose otherwise. Let $\hat{n} \neq 1$ be a firm that earns strictly positive profits. To earn positive profits, this firm must shroud and set an upfront price that attracts consumers with positive probability. Since such a price exists, $\hat{n}$ shrouds with probability one and, with probability one, chooses an upfront price that wins with positive probability. Let $\overline{f}_{\hat{n}}$ be the supremum of these prices. We distinguish two cases.

Case I: Firm $\hat{n}$ sets $\overline{f}_{\hat{n}}$ with positive probability. Then it is not a best response for firm 1 to set an upfront price f_1 above $\overline{f}_{\hat{n}} + \Delta v - \Delta v'$ because with such upfront prices firm 1 earns zero profits while it earns positive profits when offering a contract $(\underline{f}, \bar{a})$. Thus, firm 1 sets upfront prices $f_1 \leq \overline{f}_{\hat{n}} + \Delta v - \Delta v'$, contradicting the fact that firm $\hat{n}$ wins with positive probability when setting $\overline{f}_{\hat{n}}$.

Case II: Firm $\hat{n}$ sets $\overline{f}_{\hat{n}}$ with zero probability. Hence, for every $\epsilon > 0$, firm $\hat{n}$ sets base prices in the interval $\left(\overline{f}_{\hat{n}} - \epsilon, \overline{f}_{\hat{n}}\right)$ with positive probability; and this probability goes to zero as $\epsilon \to 0$. Let $\gamma \leq 1$ be the probability that all firms shroud. Firm $\hat{n}$ earning positive profits implies that $\gamma > 0$. Then, firm 1 earns equilibrium profits of at least $\gamma s_1\left(\underline{f} + \bar{a} - c\right) > 0$, which it can ensure by shrouding and offering the contract $\left(\underline{f}, \bar{a}\right)$. Since as $\epsilon \to 0$ firm 1's profits go to zero when setting an upfront price weakly above $\overline{f}_{\hat{n}} - \epsilon + \Delta v - \Delta v'$, there exists an $\overline{\epsilon} > 0$ such that firm 1 earns lower profits when setting an upfront price weakly above $\overline{f}_{\hat{n}} - \overline{\epsilon} + \Delta v - \Delta v'$ than when shrouding and offering the contract $\left(\underline{f}, \bar{a}\right)$. Hence, firm 1 sets base prices at or below $\overline{f}_{\hat{n}} - \overline{\epsilon} + \Delta v - \Delta v'$, contradicting the fact that firm $\hat{n}$ wins with positive probability when setting prices in the interval $\left(\overline{f}_{\hat{n}} - \overline{\epsilon}, \overline{\epsilon}\right)$.

Finally, we show that in any equilibrium in which all firms $n \neq 1$ shroud with probability 1, firm 1 shrouds and offers the contract $\left(\underline{f} + \Delta v - \Delta v', \bar{a}\right)$ with probability one; hence, consumers weakly prefer firm 1, and firm 1 gets the entire market. If firms $n \neq 1$ shroud with positive probability, firm 1 can earn positive profits by shrouding and offering the above contract. Hence, firm 1 shrouds with probability one. Furthermore, since firm 1 makes positive profits only conditional on all rivals shrouding, it must set $a_1 = \bar{a}$ in any such equilibrium. Conditional on all firms shrouding, firm 1 attracts all consumers with probability 1 when setting $\underline{f} + \Delta v - \Delta v'$; hence, firm 1 does not charge a lower upfront price in such an equilibrium. Finally, firm 1 cannot charge strictly more than $\underline{f} + \Delta v - \Delta v'$ with positive probability because otherwise some firm $n \neq 1$ could make positive profits when shrouding and offering the contract $\left(\underline{f}, \bar{a}\right)$, which contradicts the fact that all firms $n \neq 1$ earn zero profits. ■

PROOF OF PROPOSITION 4:

Absent innovation, firm 1 earns $s_1(\underline{f} + \bar{a} - c)$ in the deceptive equilibrium we select. The proof of Part (i) of Proposition 1, which applies unaltered when there is a binding price floor given our tie-breaking assumptions, establishes that firm 1 earns $\Delta v - \Delta v'$ in the pricing subgame following a value-increasing innovation. Thus, $I_v^* = \Delta v - \Delta v' - s_1(\underline{f} + \bar{a} - c)$. Since $\underline{f} + \bar{a} > c$, this cutoff is strictly less than that in Proposition 3. ■

PROOF OF PROPOSITION 5:

Standard Bertrand arguments analogous to those in the proof of Lemma 2 imply that in any Nash equilibrium, all firms charge $f_n = \underline{f}$.

Let $\Delta a = \Delta v = \Delta$. We have

$$I_v^* = s_n\big(D(\underline{f} - \Delta) - D(\underline{f})\big)(\underline{f} + \bar{a} - c) = s_n \Delta D'(f')(\underline{f} + \bar{a} - c)$$

for some $f' \in [\underline{f} - \Delta, \underline{f}]$ (where in the last step we have used the mean value theorem). We also have

$$I_a^* = s_n \Delta D(\underline{f}).$$

Hence, we want to show that

$$D(\underline{f}) > D'(f')(\underline{f} + \bar{a} - c).$$

Since $\underline{f} < f^m$, this inequality holds if $f' = \underline{f}$. Hence, by the continuity of $D'(f)$, it also holds for Δ sufficiently small. ■

REFERENCES

Armstrong, Mark, and John Vickers. 2012. "Consumer Protection and Contingent Charges." *Journal of Economic Literature* 50 (2): 477–93.

Ausubel, Lawrence M. 1991. "The Failure of Competition in the Credit Card Market." *American Economic Review* 81 (1): 50–81.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence." *Quarterly Journal of Economics* 119 (2): 353–402.

Eliaz, Kfir, and Ran Spiegler. 2006. "Contracting with Diversely Naive Agents." *Review of Economic Studies* 73 (3): 689–714.

Ellison, Glenn. 2005. "A Model of Add-On Pricing." *Quarterly Journal of Economics* 120 (2): 585–637.

Evans, David S., and Richard Schmalensee. 2005. *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. Cambridge: MIT Press.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121 (2): 505–40.

Grubb, Michael D. 2009. "Selling to Overconfident Consumers." *American Economic Review* 99 (5): 1770–1805.

Grubb, Michael D. 2015. "Consumer Inattention and Bill-Shock Regulation." *Review of Economic Studies* 82 (1): 219–57.

Hall, Robert E. 1997. "The Inkjet Aftermarket: An Economic Analysis." http://web.stanford.edu/~rehall/Inkjet%20Aftermarket%201997.pdf.

Heidhues, Paul, and Botond Kőszegi. 2010. "Exploiting Naivete about Self-Control in the Credit Market." *American Economic Review* 100 (5): 2279–2303.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2012a. "Deception and Consumer Protection in Competitive Markets." In *The Pros and Cons of Consumer Protection*, edited by Dan Sjöblom, 44–76. Stockholm: Swedish Competition Authority.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2012b. "Inferior Products and Profitable Deception." http://federation.ens.fr/ydepot/semin/texte1213/BOT2013INF.pdf.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2014. "Inferior Products and Profitable Deception." http://www.personal.ceu.hu/staff/Botond_Koszegi/inferior_products.pdf.

Ko, K. Jeremy. 2012. "Disclosure and Price Regulation in a Market with Potentially Shrouded Costs." Unpublished.

Lal, Rajiv, and Carmen Matutes. 1994. "Retail Pricing and Advertising Strategies." *Journal of Business* 67 (3): 345–70.

Piccione, Michele, and Ran Spiegler. 2012. "Price Competition under Limited Comparability." *Quarterly Journal of Economics* 127 (1): 97–135.

Reinganum, Jennifer F. 1989. "The Timing of Innovation: Research, Development, and Diffusion." In *Handbook of Industrial Organization*, Vol. 1, edited by Richard Schmalensee and Robert D. Willig, 849–908. Amsterdam: North-Holland.

Spiegler, Ran. 2006a. "Competition over agents with boundedly rational expectations." *Theoretical Economics* 1 (2): 207–31.

Spiegler, Ran. 2006b. "The Market for Quacks." *Review of Economic Studies* 73 (4): 1113–31.

Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization*. Oxford: Oxford University Press.

Stango, Victor, and Jonathan Zinman. Forthcoming. "Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees." *Review of Financial Studies*.