Exploitative Innovation*

Paul Heidhues
ESMT, CEPR and CESifo

Botond Kőszegi
Central European University and UC Berkeley

Takeshi Murooka
UC Berkeley

December 3, 2012

Abstract

We analyze innovation incentives in a simple model of a competitive retail market with naive consumers. Firms selling perfect substitutes play a game consisting of an innovation stage and a pricing stage. At the pricing stage, firms simultaneously set a transparent "up-front price" and an "additional price," and decide whether to shroud the additional price from naive consumers. To capture especially financial products such as banking services, credit cards, and mutual funds, we allow for a floor on the product's up-front price. At the preceding innovation stage, a firm can invest either in increasing the product's value *(value-increasing innovation)* or in increasing the maximum additional price *(exploitative innovation)*. We show that if the price floor is not binding, the incentive for either kind of innovation equal the "appropriable part" of the innovation, implying similar incentives for exploitative and value-increasing innovations. If the price floor is binding, however, innovation incentives are often stronger for exploitative than for value-increasing innovations. Because learning ways to charge higher additional prices increases the profits from shrouding and thereby lowers the motive to unshroud, a firm may have strong incentives to make appropriable exploitative innovations, and *even stronger* incentives to make non-appropriable exploitative innovations. In contrast, the incentive to make non-appropriable value-increasing innovations is zero or negative, and even the incentive to make appropriable value-increasing innovations is strong only if the product is socially wasteful. These results help explain why firms in the financial industry have been willing to make innovations others could easily copy, and why these innovations often seem to have included exploitative features. **JEL Codes: D14, D18, D21**

*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," and can be found at http://elsa.berkeley.edu/~botond/inferiorproducts.pdf. 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.

1 Introduction

A growing theoretical literature in behavioral economics investigates how firms use hidden fees—e.g., overdraft fees for bank accounts, late fees and high interest payments for credit cards, and charges triggered by complex reset rules in mortgages—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 consumers value. Section 2 introduces our model, which consists of a simultaneous-move price-competition stage modeled after Gabaix and Laibson (2006) and Heidhues, Kőszegi and Murooka (2012), and a preceding innovation stage. At the price-competition stage, firms selling perfect substitutes each set a transparent up-front price as well as an additional price, and unless at least one firm decides to (costlessly) unshroud additional prices, naive consumers ignore these prices when making purchase decisions. To capture the notion that in some markets, such as banking services, credit cards, and mutual funds, firms cannot return all profits from later charges by lowering initial charges—even if they can dissipate those profits in other ways—we deviate from most existing work and posit that there is a floor on the up-front price.[1] We assume that whenever a deceptive equilibrium—wherein all firms shroud additional prices—exists at the pricing stage, firms play that continuation equilibrium.[2]

[1] As we discuss in Section 2, in Heidhues, Kőszegi and Murooka (2011) we provide a microfoundation for the price floor based on the presence of "arbitrageurs" who would take advantage of overly low prices. This microfoundation is an extreme variant of Ellison's (2005) insight that firms may be reluctant to cut initial prices because these cuts disproportionately attract less profitable consumers. While we assume that firms may not be able to return all profits by lowering the up-front price, however, our model is consistent with the possibility that firms dissipate profits in other ways—such as through entry costs or advertising—and hence earn zero net profits. We discuss this issue in more detail in the conclusion.

[2] 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 equilibrium from

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 socially wasteful exploitative innovation—increasing the maximum additional price—or in socially often beneficial value-increasing innovation—increasing the product's value—and other firms observe its innovation decision. We consider both appropriable innovations (which other firms cannot copy) and non-appropriable innovations (which other firms can fully copy), as well as in-between cases.

In Section 3, we characterize innovation incentives when the price floor is not binding—a condition we argue holds for some commonly invoked examples of deceptive products, including hotel rooms and printers—and find that they are, identically for exploitative and value-increasing innovations, 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 incentives 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 4, we characterize innovation incentives when the price floor is binding—a condition we argue holds for a number of consumer financial products, including credit cards and banking services—and find that they are now often stronger for exploitative than for value-increasing innovation. Our results can be understood in terms of firm 1's incentive to maintain a deceptive equilibrium in the face of the threat that a competitor—unable to compete on the up-front price—prefers to unshroud and compete on the additional price. This deviation allows a competitor to capture the entire market, but because consumers who learn of the additional prices may not buy the product, it may require a discrete cut in the total price. A profitable deceptive continuation

the perspective of the firms, and in some settings it does so strictly.

equilibrium, therefore, exists if and only if each firm's profit when sharing deceived consumers at a high profit margin is greater than its profit when unshrouding and capturing all consumers at a low profit margin.

The above logic implies that so long as a profitable deceptive continuation equilibrium exists, firm 1 has strong incentives for both appropriable and non-appropriable exploitative innovation. Since in a deceptive equilibrium the floor on the up-front price is binding, firm 1 always benefits from an exploitative innovation by being able to increase its margin. Furthermore, because acquiring the exploitative innovation increases competitors' profits from shrouding and thereby lowers their motive to unshroud, the innovation may enable profitable deception in the industry. In this case, firm 1 *prefers* competitors to acquire the exploitative innovation—so that it prefers non-appropriable to appropriable innovations—and its willingness to pay is equal to its full post-innovation profits. In contrast, firm 1 can never benefit from a non-appropriable value-increasing innovation, and because such an innovation increases competitors' motive to unshroud by raising the profits from selling a transparent product, a firm may even be willing to pay to *avoid* the innovation.

To complete our analysis, we explore the incentives for fully or partially appropriable value-increasing innovations, showing that even these are strong only for products that should not survive in the market in the first place. Note that in a socially valuable industry—where consumers value the product above production cost—a competitor who would otherwise attract no consumer prefers to unshroud. Hence, because an appropriable value-increasing innovation steals the consumers of other firms, in a socially valuable industry it leads to unshrouding and the loss of profits from deception. As a result, the incentive for such innovation in a socially valuable industry is weak, and is negative for small increases in value. In a socially wasteful industry—where consumers value the product below production cost—in contrast, a firm that unshrouds cannot go on to profitably sell its product, so no firm has an incentive to unshroud. With unshrouding not a concern, firm 1 prefers to steal competitors' consumers to increase profits. As a result, its innovation incentive is strong, and is non-trivial even for vanishingly small increases in value.

In many situations, therefore, our model implies substantial incentives for exploitative innovation. From a welfare perspective, all of this spending is of course 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 has 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.[3] 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 prevalent theme in the classical literature, that firms selling substitutes often underinvest in non-appropriable innovations (e.g., Reinganum 1989).

2 Basic Model

In this section, 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 et al. (2012), 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 up-front 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 $\overline{a}_n > 0$, 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

[3] 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 (2010), Grubb (2012), and Ko (2012), and Spiegler (2011) for a textbook treatment.

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 $f_n + a_n$.[4]

Crucially, we deviate from much of the literature and impose that firms face a floor on the up-front price: $f_n \geq \underline{f}$. We do not impose a floor on the total price. In Heidhues et al. (2011), we provide one microfoundation for the floor on the up-front price based on the threat of "arbitrageurs" who can avoid the additional price because they are not interested in using the product itself, and who would therefore bankrupt a firm that cuts its up-front price below their entry threshold (e.g., \$0).[5] This microfoundation is an extreme variant of Ellison's (2005) insight (developed in the context of add-on pricing) that firms may be reluctant to cut initial prices because these cuts disproportionately attract less profitable consumers. Armstrong and Vickers (2012), Grubb (2012), and Ko (2012) also analyze models with variants of our price-floor assumption. While we assume that firms may not be able to return all profits by lowering the up-front price, however, our model is consistent with the possibility that firms dissipate profits in other ways—such as through entry costs or advertisements—and hence earn zero net profits. We discuss this issue in more detail in the conclusion.

Each firm's cost of providing the product is $c > 0$. We assume that $v_n + \overline{a}_n > c$ for all n; a firm with $v_n + \overline{a}_n < c$ cannot profitably sell its product, so without loss of generality we can think of it as not participating in the market. We also impose 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 up-front price at the floor. Finally, we make two simple tie-breaking assumptions. First, consumers go to a highest-quality firm when indifferent.[6]

[4] The above framework incorporates two unconventional assumptions: that there is an additional price component and that firms make an unshrouding decision regarding this component. The assumption that there is a price component naive consumers may ignore and cannot avoid is a reduced-form version of many assumptions that have appeared in the behavioral industrial-organization literature (DellaVigna and Malmendier 2004, Gabaix and Laibson 2006, for example). Our qualitative results remain unchanged if (as in most previous theories) naive consumers can partially but not fully avoid the additional price. Our assumption that firms can unshroud additional prices to all consumers is an extreme way of capturing the notion that firms may use education to attract consumers from competitors. The logic of our qualitative results does not seem to depend on firms being able to educate all consumers.

[5] As a simple illustration in a specific case, consider the finding of Hackethal, Inderst and Meyer (2010) that German bank revenues from security transactions amount to €2,560 per customer per year (2.43% of mean portfolio value). If a bank handed out such sums ex ante—even if it did so net of account maintenance costs—many individuals would sign up for (and then not use) bank accounts just to get the handouts.

[6] 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.

Second, if all firms shroud and a subset of firms with the same quality choose an up-front 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. Firms start from a symmetric position in v_n and $\overline{a}_n$, with $v_n = v$ and $\overline{a}_n = \overline{a}$ for all n. 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. 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 $0 \leq \Delta v' \leq \Delta v$.[7]

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 Nash 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. 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 then the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is.[8] In addition, in most of our settings a deceptive

[7] 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.

[8] The proof of Proposition 5 in Heidhues et al. (2012) establishes this claim for any situation in which consumers value firms' products equally and the price floor is binding, a condition that applies to Propositions 2 and 3 below. For Part II of Proposition 1 and Proposition 4 below, we establish the claim in Appendix B. In Proposition 5 below, we use our equilibrium refinement only for the subgame following no innovation, and in this case consumers value

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.

3 Non-Binding Price Floor

We first consider innovation incentives when the floor on the up-front price is not binding, supposing throughout this section that $\underline{f} \leq c - (\overline{a} + \Delta a)$. This condition seems to hold for some commonly invoked examples of deceptive products, such as printers and hotel rooms (Hall 1997, Gabaix and Laibson 2006, for instance). As an illustrative example, the marginal cost of a hotel room (c) is likely to be non-trivial, and the additional amount a hotel can extract from the minibar, telephone, and other add-on services ($\overline{a} + \Delta a$) is limited. Hence, with a price floor of around \$0 the above inequality is satisfied. Consistent with this view, the up-front (base) price of hotels and printers is typically well above \$0.

Proposition 1 identifies innovation incentives in this case:

Proposition 1. *Suppose* $\underline{f} \leq c - (\overline{a} + \Delta a)$ *for all* n. *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 that of the other firms 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 up-front

products equally and the price floor is binding. In Part I of Proposition 1, we do not use the refinement.

price to \$100-\$50=\$50. Absent the innovation, therefore, firm 1 cannot sell its product above an up-front price of \$50, so it earns zero net profits. If it innovates, however, firm 1 can charge an up-front price slightly below \$60 and attract all consumers, generating total revenue of nearly \$110 with the additional price included. 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 up-front price to \$40. Even so, firm 1 can offer a slightly lower up-front price, and again earn a revenue of nearly \$110 with its higher additional price included. And because other firms are charging a total price equal to marginal cost, no firm can profitably unshroud and attract consumers. 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.

4 Binding Price Floor

In this section, we analyze innovation incentives when the price floor is binding: $\underline{f} > c - \overline{a}$. This case describes a number of consumer financial products, including credit cards, bank accounts, and actively-managed mutual funds. For instance, the floor on the up-front price of a credit card ($\underline{f}$) is not much below \$0, and the marginal cost of setting up a credit-card account to a consumer (c) is also quite low. At the same time, credit-card companies make substantial amounts in hidden fees ($\overline{a}$ is large), so that the above inequality is easily satisfied.[9,10]

To analyze innovation incentives, we first identify conditions under which a deceptive equilibrium exists in the pricing subgame, assuming that the maximum additional price firm n can charge is $\overline{a}_n$ and the value of its product is v. The analysis mirrors that in Heidhues et al. (2012). Note that if additional prices remain shrouded, all firms set their maximum additional price $\overline{a}_n$. Then, since firms are making positive profits and hence have an incentive to attract consumers, they bid down the up-front price to $\underline{f}$. With consumers being indifferent between firms, firm n gets market share s_n and therefore earns a profit of $s_n(\underline{f} + \overline{a}_n - c)$. 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 can make profits of at most $v - c$ by unshrouding and capturing the entire market. Hence, unshrouding is unprofitable for firm n if

$$s_n(\underline{f} + \overline{a}_n - c) \geq v - c, \tag{SC}$$

and a deceptive equilibrium exists if this "Shrouding Condition" holds for all n. The following lemma summarizes this condition, and also identifies what happens when Condition (SC) is violated for some firm:

[9] If credit-card companies handed out substantial sums to new consumers (i.e., if the up-front price was substantially negative), then many deal-seeking consumers would likely get—and then not use—these cards just to earn the handouts. The threat of losing money on these "arbitrageur" consumers imposes a floor on the up-front price that is likely not much below \$0. 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.

[10] A less clear-cut example is mortgages. In this market, the up-front price can correspond to initial monthly payments and the additional price to future monthly payments, prepayment penalties and other fees. Then, it is clear that arbitrageurs of the type above do not impose a floor on the up-front price. But if—similarly to Ellison (2005)— cutting the initial monthly payments to very low levels would attract primarily risky borrowers, firms might not compete too much on this up-front price, imposing something akin to a price floor.

Lemma 1 (Equilibrium in the Pricing Subgame)**.** *Suppose* $\underline{f} > c - \overline{a}_n$ *for all* n*. If Inequality* (SC) *holds for all* n*, a deceptive continuation equilibrium exists. In any deceptive continuation equilibrium,* $f_n = \underline{f}$ *and* $a_n = \overline{a}_n$ *for all* n*, and firms earn positive profits. If Inequality* (SC) *is violated for some* n*, 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 up-front terms. But the price floor prevents firms from competing away all profits from the additional price by lowering the up-front price. Second, since firms cannot compete for consumers by cutting their up-front price, there is a pressure for competition to shift to the additional price—but because competition in the additional price requires unshrouding, it is an imperfect substitute for competition in the up-front 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 et al. (2012), 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 consumer 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 also show that the relevant notion of social wastefulness is relative to the best alternative: whether or not the product yields positive social surplus, if an alternative product with higher social surplus is available, there is typically a deceptive equilibrium in which naive consumers buy the inferior product and firms earn positive profits from it.

Our key results derive from considering how an innovation by firm 1 affects firms' willingness to go along with deceiving consumers, as captured by the Shrouding Condition (SC). Proposition 2 states our results for non-appropriable innovations, showing that they are much stronger for exploitative than for value-increasing innovations:

Proposition 2 (Non-Appropriable Innovations)**.** *Suppose* $\underline{f} > c - \overline{a}$*. Then,*

I. (Exploitative.) Suppose $\Delta a = \Delta a'$*. If all firms satisfy the Shrouding Condition for* $\overline{a}_n = \overline{a} + \Delta a$*, then* $I_a^* \geq s_1 \Delta a$*. If in addition some firm* n *does not satisfy the Shrouding Condition for* $\overline{a}_n = \overline{a}$*, then* $I_a^* = s_1(\underline{f} + \overline{a} + \Delta a - c) > s_1 \Delta a$*.*

II. (Value-Increasing.) Suppose $\Delta v = \Delta v'$*. Then,* $I_v^* \leq 0$*.*

Part I of Proposition 2 says that if firms can maintain a deceptive equilibrium following the innovation, firm 1 is willing to spend resources on—socially clearly wasteful—non-appropriable exploitative innovation. In this situation, increasing the additional price from $\overline{a}$ to $\overline{a} + \Delta a$ cannot lead to a decrease in the up-front price, so the innovation increases firm 1's profits by at least $s_1 \Delta a$. Going further, if in addition firms cannot maintain a deceptive equilibrium without the innovation, the innovation increases firm 1's profits even more by enabling profitable deception in the industry. In this case, firm 1's willingness to pay for the exploitative innovation is equal to its full post-innovation profits—a potentially huge incentive to innovate.

Part II of Proposition 2 says that in contrast to socially-wasteful exploitative innovation, firm 1 is not willing to spend on—socially often beneficial—non-appropriable value-increasing innovation. Because such innovation can increase neither one's market share nor one's markup, firm 1 has no incentive to invest in it. Furthermore, it may be the case that firms can maintain a deceptive equilibrium without but not with the 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.

To sharpen the intuitions from Proposition 2 on a firm's incentive to invest in exploitative innovation, we compare the incentives for non-appropriable and appropriable innovations:

Proposition 3 (Appropriability of Exploitative Innovation)**.** *Suppose* $\underline{f} > c - \overline{a}$*. Then,* I_a^* *is weakly greater if the innovation is non-appropriable* ($\Delta a = \Delta a'$) *than if it is only partially copyable* ($\Delta a \geq \Delta a'$)*, and it is strictly greater if the Shrouding Condition holds for all* n *when* $\overline{a}_n = \overline{a} + \Delta a$ *but fails for some* $n \neq 1$ *when* $\overline{a}_n = \overline{a} + \Delta a'$*.*

Firm 1 has a weakly *greater* incentive to engage in exploitative innovation if other firms can copy its innovation than if they cannot—and, equivalently, firm 1 weakly prefers others to obtain its innovation. As above, an increase in a competitor's additional price does not lead to greater competition in the up-front price, so it never lowers firm 1's profits. Moreover, 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.[11]

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. As has been argued by many researchers, consumers likely underappreciate a number of future fees and other payments associated with credit cards, bank accounts, and non-traditional mortgages, and firms exacerbate these mispredictions with carefully designed contract features—such as teaser rates, high fees for certain patterns of product use, and difficult-to-understand payment schedules involving large future payments—whose main purpose is likely to hide products' total price from consumers. These exploitative contract innovations seem easy to copy, and in fact have been quickly copied by competitors.

Furthermore, not only are the above types of innovations easy to copy, in some instances firms seem—consistent with Proposition 3—positively willing to share them with each other. For example, Argus is an information-exchange service that collects individual-level account data from

[11] An important caveat to Proposition 3 is that while firm 1 is willing to share an exploitative innovation with competitors already in the marketplace, it often prefers not to share the same innovation with potential entrants. The ability to charge higher additional prices can make participation in the market more attractive, inducing additional entry and thereby reducing the innovator's market share.

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."[12] For a participating issuer, any innovation is essentially a non-appropriable innovation. Proposition 2 explains why a participating firm makes innovations, and although there may be other considerations, Proposition 3 provides a strategic reason for why a firm that is interested in developing new exploitative practices is willing to join Argus in the first place.

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

Proposition 4 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 non-trivial 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 up-front 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:

[12] See `http://www.argusinformation.com/eng/our-services/syndicated-studies/credit-card-payment-study/us-credit-card-payments-study/`.

Proposition 5 (Value-Increasing Innovation in Socially Valuable Industries)**.** *Suppose* $\underline{f} > c - \overline{a}$*,* $v > c$*,* $\Delta v > \Delta v'$*, and (SC) holds for all* n *when* $\overline{a}_n = \overline{a}$*. Then,* $I_v^* = [\Delta v - \Delta v'] - s_1(\underline{f} + \overline{a} - c)$*.*

Proposition 5 implies that firm 1's willingness to pay for partially appropriable value-increasing innovations in a socially valuable deceptive industry is quite small: it is lower than in the corresponding classical setting (where it would equal $\Delta v - \Delta v'$), it is lower than in a socially wasteful industry with the same relative product values, and it is negative for non-substantial improvements. In this industry, any partially appropriable innovation must lead to unshrouding; otherwise, firm 1 would be able to price competitors out of the market while setting a high total price, and competitors could unshroud and profitably undercut this total price. Hence, innovation leads firm 1 to lose its positive profits from deception. This loss dampens firm 1's incentive to innovate, and for small improvements—which generate only a small relative advantage—the incentive is negative.

5 Concluding Remarks

A key implication of our model is that the floor may prevent firms from competing away all profits from deception by lowering the up-front price. Nevertheless, our theory is consistent with a number of other ways in which firms can dissipate the profits from deception. We can distinguish two different types of opportunities to dissipate profits. If these opportunities arise *before* the innovation decision, they are sunk at the point at which our game starts, and our results regarding innovation incentives remain unchanged. What is more, the welfare effects of exploitative innovation are worse in this case than in our basic model above: not only do firms spend resources on a socially wasteful activity, in equilibrium they might dissipate the resulting profits in a socially worthless activity as well. For instance, firms might dissipate part or all of their profits through simple entry costs or through initial marketing or advertising—perhaps aimed to increase their market share among consumers who are indifferent at the later stage.

If opportunities to dissipate profits arise *after* the innovation stage, however, our results may be affected. In particular, if firms can offer perks at the pricing stage that are just as effective in attracting consumers as decreases in the up-front price, our theory becomes equivalent to one without a price floor. But it seems plausible to assume that non-price forms of competition are often less effective in attracting consumers, and are subject to decreasing returns. For example, while mail solicitations clearly generate some demand for credit cards, the returns to sending multiple solicitations to a consumer within a short time frame likely diminish quickly. In this case, firms do not compete away all profits from the additional price at the pricing stage. As a result, the logic of our main insights remains unchanged, albeit with an important qualification to our prediction that a firm always prefers a non-appropriable to an appropriable exploitative innovation. Specifically, a firm is still willing to spend on a non-appropriable innovation, and prefers it over an appropriable innovation if the innovation is pivotal in preventing another firm from unshrouding. But because a fraction of the extra profits from the innovation is competed away, if the innovation is not pivotal in preventing unshrouding the firm prefers to appropriate it.

Given our emphasis that exploitation requires innovation, it would be interesting to investigate the dynamics of how exploitation has appeared and spread in an industry. One possible scenario suggested by our theory is the following. The industry is initially in a non-deceptive situation (e.g. offering only 30-year fixed-rate mortgages). Then, one firm invents and starts offering a product with shroudable features (e.g. a complex adjustable-rate mortgage), and 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 line. Our theory suggests that competitors' preference is to adopt the deceptive product.

References

Armstrong, Mark and John Vickers, "Consumer Protection and Contingent Charges," MPRA Paper 37239, University Library of Munich, Germany 2012.

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

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

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

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

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

Gabaix, Xavier and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

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

____ , "Consumer Inattention and Bill-Shock Regulation," 2012. Working Paper, MIT Sloan School of Business.

Hackethal, Andreas, Roman Inderst, and Steffen Meyer, "Trading on Advice," 2010. Working Paper, University of Frankfurt.

Hall, Robert E., "The Inkjet Aftermarket: An Economic Analysis," 1997. Unpublished Manuscript, Stanford University.

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

____ , ____ , and Takeshi Murooka, "Deception and Consumer Protection in Competitive Markets," 2011. Anthology on the Pros and Cons of Consumer Protection; `http://elsa.berkeley.edu/~botond/stockholm.pdf`.

____ , ____ , and ____ , "Inferior Products and Profitable Deception," 2012. Working Paper.

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

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

Piccione, Michele and Ran Spiegler, "Price Competition under Limited Comparability," 2010. Working paper, LSE and UCL.

Reinganum, Jennifer F., "The Timing of Innovation: Research, Development, and Diffusion," in Richard Schmalensee and Robert D. Willig, eds., *Handbook of Industrial Organization*, Vol. 1, Amsterdam; Oxford and Tokyo: North-Holland, 1989, pp. 849–908.

Spiegler, Ran, "Competition over Agents with Boundedly Rational Expectations," *Theoretical Economics*, 2006, *1* (2), 207–231.

____ , "The Market for Quacks," *Review of Economic Studies*, 2006, *73* (4), 1113–1131.

____ , *Bounded Rationality and Industrial Organization*, Oxford University Press, 2011.

Appendix A: Proofs

Proof of Lemma 1. See proof of Proposition 3 in Heidhues et al. (2012). □

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 $\overline{a}$, firm 1 sets an up-front price of $c - \overline{a} + \Delta v - \Delta v'$, all other firms set an up-front price of $c - \overline{a}$, and firm 1 gets the entire market. Firm 1 earns a profit of $\Delta v - \Delta v'$.

We next argue that firm 1 earns at least $\Delta v - \Delta v'$ in any equilibrium. Recall that by assumption no firm prices below marginal cost. Thus, if firm 1 unshrouds and and charges a total price of $c + \Delta v - \Delta v'$ all consumers weakly prefer the product of firm 1 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. 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; 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 it must set $a_k = \overline{a}$ to maximize profits, and hence it offers a contract $(\hat{t} + \Delta v' - \overline{a}, \overline{a})$ with positive probability.

For this contract to have a positive market share, firm 1 must shroud and set base prices above $\hat{t} + \Delta v - \overline{a}$ with positive probability, and for its quality-adjusted price to be below $\hat{t}$ firm 1 at the same time must set the additional price below $\overline{a} - \Delta v$. But this is not a best response: firm 1 could keep the total price distribution fixed but always charge the maximal additional price $\overline{a}$. Then its base price would always lie (weakly) below $\hat{t} + \Delta v - \overline{a}$, and firm 1 could 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 cannot do so when unshrouding as it would have zero market share. And firm 1 cannot charge this price with positive probability when shrouding: if it did, firm k would have to charge base price weakly above $\hat{t} + \Delta v' - \overline{a}$ when shrouding for firm 1 to have positive market share. But firm k could keep its total price distribution fixed and always charge the maximal additional price $\overline{a}$, thereby ensuring that its base price is below $\hat{t} + \Delta v' - \overline{a}$, increasing its market share while holding the total price fixed. We conclude that neither 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, so that 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, then, that firm 1 shrouds on this interval, but its market share does not approach zero. Then firm k must with positive probability shroud and set up-front prices at or above $\hat{t} + \Delta v' - \overline{a}$ while setting total prices strictly below $\hat{t} + \Delta v'$. Firm k in this case could keep its total price distribution fixed, and always charge the maximal additional price $\overline{a}$, increasing its market share for a set of total prices it charges 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 $(\hat{t} - \epsilon, \hat{t})$ is analogous. We conclude that firm 1 earns its relative advantage in every equilibrium.

Part II. It is easy to check that the following constitutes an equilibrium in the pricing subgame:

all firms shroud and set $f_n = c-(\overline{a}+\Delta a')$ and their maximum additional price, firm 1 gets the entire market, and earns a profit of $\Delta a - \Delta a'$. We show that profits are the same in any continuation equilibrium in which firms shroud with probability one, immediately implying the proposition. Clearly, in any continuation equilibrium in which firms shroud with probability one, all firms set the maximum additional price. Then, the game is equivalent to a Bertrand-competition game in which firm 1 has cost $c-(\overline{a}+\Delta a)$ and all other firms have cost $c-(\overline{a}+\Delta a')$. It is well-known that in any equilibrium of this game when no firm charges below marginal cost, firm 1 earns a profit equal to $\Delta a - \Delta a'$. □

Proof of Proposition 2. We first prove Case I. In the subgame following an innovation by firm 1, the Shrouding Condition holds for all firms by assumption, and thus firm 1 earns $s_1(\underline{f}+\overline{a}+\Delta a-c)$ in this case. In the subgame in which firm 1 did not innovate, firm 1 earns $s_1(\underline{f}+\overline{a}-c)$ 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 $s_1(\underline{f}+\overline{a}+\Delta a-c)$, which is strictly greater than $s_1\Delta a$ because $\underline{f}+\overline{a}>c$.

We now prove Case II. 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 $s_1(\underline{f}+\overline{a}-c)$ 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. Firm 1 earns $s_1(\underline{f}+\overline{a}+\Delta a-c)$ in the subgame following its innovation if the Shrouding Condition holds for all firms and zero profits otherwise. An increase in $\Delta a'$ for some $n \neq 1$ increases the left-hand-side of the Shrouding Condition and hence relaxes it; thus it either does not affect firm 1's profits or—if it makes the Shrouding Condition hold for some firm $n \neq 1$ for which it does not hold otherwise—strictly increases firm 1's profits. □

Proof of Proposition 4.

We 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 that firms play this equilibrium in the pricing subgame whenever it exists) firm 1 therefore earns $s_1(\underline{f} + \overline{a} - c)$.

Now 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 $(\underline{f} + \Delta v - \Delta v', \overline{a})$ with probability one, and all firms $n \neq 1$ offer the contract $(\underline{f}, \overline{a})$; 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 a firm $n \neq 1$ that yields positive profits. If firm 1 deviates and unshrouds or shrouds and sets a higher base price, it earns zero profits. And since it has a market share of one, firm 1 cannot benefit from lowering its base 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 1, firm 1 charges prices $\underline{f} + \Delta v - \Delta v', \overline{a}$ 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 $s_1(\underline{f} + \overline{a} - c)$ 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 up-front price that attracts consumers with positive probability. Since such a price exists, $\hat{n}$ shrouds with probability one and, with probability one, chooses an up-front 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 up-front price f_1 above $\overline{f}_{\hat{n}} + \Delta v - \Delta v'$ because with such base prices firm 1 earns zero profits while it earns positive profits when offering a contract $(\underline{f}, \overline{a})$. Thus, firm 1 sets base 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 $(\overline{f}_{\hat{n}} - \epsilon, \overline{f}_{\hat{n}})$ with positive probability; and this probability goes to zero as $\epsilon \to 0$. Let $\gamma \leq 1$ be the probability that all firms shroud. That firm $\hat{n}$ earns positive profits implies that $\gamma > 0$. Then, firm 1 earns equilibrium profits of at least $\gamma s_1(\underline{f} + \overline{a} - c) > 0$, which it can ensure by shrouding and offering the contract $(\underline{f}, \overline{a})$. Since as $\epsilon \to 0$ firm 1's profits go to zero when setting a base price at or 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 a base price at or above $\overline{f}_{\hat{n}} - \overline{\epsilon} + \Delta v - \Delta v'$ than when shrouding and offering the contract $(\underline{f}, \overline{a})$. 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 $(\overline{f}_{\hat{n}} - \overline{\epsilon}, \overline{\epsilon})$.

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 $(\underline{f} + \Delta v - \Delta v', \overline{a})$ with probability one; hence, consumers weakly prefer fir m 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 $a_1 = \overline{a}$ in any such equilibrium. Since conditional on all firms shrouding, firm 1 attracts all consumers with probability 1 when setting $\underline{f} + \Delta v - \Delta v'$, firm 1 does not charge a lower base 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 $(\underline{f}, \overline{a})$, which contradicts the fact that all firms $n \neq 1$ earn zero profits. □

Proof of Proposition 5. Absent innovation, firm 1 earns $s_1(\underline{f} + \overline{a} - c)$ in the deceptive equilibrium we select. The proof of Part I of Proposition 1, which applies unaltered when there is a binding price floor, establishes that firm 1 earns $\Delta v - \Delta v'$ in the pricing subgame following a value-increasing innovation.Thus, $I_v^* = \Delta v - \Delta v' - s_1(\underline{f} + \overline{a} - c)$. Since $\underline{f} + \overline{a} > c$, this cut-off is strictly less than that in Proposition 4. □

Appendix B: Uniqueness of Deceptive Equilibrium with Unshrouding Cost

Part II of Proposition 1. We 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.

(i): *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.

(ii): *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—i.e. chose the same distribution over total prices—an unshrouding firm follows conditional on unshrouding, earning 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 $f_1 = c - \overline{a} - \Delta a' - \epsilon$, $a_1 = \overline{a} + \Delta a$ 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.

(iii) *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 up-front 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 show that $j \neq 1$. Suppose, toward a contradiction, that $j = 1$. Consider a deviation by firm 1 in which it shrouds and sets $\hat{f}' = \hat{f} - (\overline{a} + \Delta a - \hat{a}) - \frac{\eta}{2}, \hat{a}' = \overline{a} + \Delta a$. 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 up-front prices set by firms other than j conditional on shrouding. Then, for firm j to earn its 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 π_k be firm k's equilibrium expected profit. By an argument structurally identical to that in the first paragraph of step (iii), for $\tilde{f}''$ sufficiently close to $\tilde{f}$ firm k cannot earn a profit of π_k when shrouding occurs, so that it 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 sufficiently small. This deviation does not affect profits if unshrouding occurs, and increases profits if shrouding occurs, a contradiction. □

Proposition 4. Obvious: an unshrouding firm makes zero gross profits as it cannot profitably sell the product, so that it makes negative net profits. □