Review of Economic Studies (2017) **84**, 323–356 doi:10.1093/restud/rdw037

Advance access publication 6 August 2016

Inferior Products and Profitable Deception*

PAUL HEIDHUES
ESMT

BOTOND KŐSZEGI
Central European University

and

TAKESHI MUROOKA
University of Munich

First version received October 2012*; final version accepted May* 2016 (*Eds.*)

We analyse conditions facilitating profitable deception in a simple model of a competitive retail market. Firms selling homogenous products set anticipated prices that consumers understand and additional prices that naive consumers ignore unless revealed to them by a firm, where we assume that there is a binding floor on the anticipated prices. Our main results establish that "bad" products (those with lower social surplus than an alternative) tend to be more reliably profitable than "good" products. Specifically, (1) in a market with a single socially valuable product and sufficiently many firms, a deceptive equilibrium—in which firms hide additional prices—does not exist and firms make zero profits. But perversely, (2) if the product is socially wasteful, then a profitable deceptive equilibrium always exists. Furthermore, (3) in a market with multiple products, since a superior product both diverts sophisticated consumers and renders an inferior product socially wasteful in comparison, it guarantees that firms can profitably sell the *inferior* product by deceiving consumers. We apply our framework to the mutual fund and credit card markets, arguing that it explains a number of empirical findings regarding these industries.

Key words: Deceptive product, Inferior product, Naivete, Profitable deception, Mutual funds, Credit cards

JEL Codes: D14, D18, D21

1. INTRODUCTION

In this article, we investigate circumstances under which firms sell products by deceiving some consumers about the products' full price, focusing (in contrast to much of the literature) on deception that leads to positive equilibrium profits in seemingly competitive industries.[1] We identify a novel, perverse aspect of profitable deception: products that generate lower social

*The setup of this article was developed in our earlier working paper "The Market for Deceptive Products" (first full version: January 2012), which we split into two papers. The companion paper is Heidhues *et al.* (2016).

1. Hidden fees have often enabled firms to reap substantial profits despite seemingly considerable competition, at least at the price-competition stage when entry and marketing costs have been paid and customer bases have been

surplus than the best alternative facilitate deception precisely because they would not survive in the market if consumers understood their full price, and, therefore, firms often make profits on exactly such bad products but not on good products. We argue that—in seeking these profits—firms disproportionately push bad products and might enter markets for bad products in potentially massive numbers, creating further economic inefficiencies.

We develop our insights in the reduced-form model of Section 2, which builds on the seminal theory of Gabaix and Laibson (2006). Firms are engaged in simultaneous-move price competition to sell a single homogenous product to naive consumers. Each firm charges a transparent anticipated price as well as an additional price, and unless at least one firm costlessly unshrouds (*i.e.*, reveals) the additional prices, consumers ignore these prices when making purchase decisions. To capture the notion that for some products 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 anticipated price.[2] We argue that a profitable deceptive equilibrium—wherein all firms shroud additional prices—is the most plausible equilibrium whenever it exists: it is then the unique equilibrium in the variant of our model in which unshrouding carries a cost (no matter how small), and all firms prefer it over an unshrouded prices and, therefore, zero-profit equilibrium. Hence, in our discussion we presume that a deceptive equilibrium is played whenever it exists. Of course, the existence of a profitable equilibrium at the pricing stage does not mean that the net profits taking all operating costs into account are large or even positive in a deceptive market. For example, entry or marketing costs can dissipate any profits from the later stage of serving consumers; indeed, these costs constitute an additional waste generated by profitable deception.

Section 3 presents our basic results. As a benchmark, we show that if the price floor is not binding, only a zero-profit deceptive equilibrium exists. If the price floor is binding, however, profitable deception may occur. If other firms shroud and the anticipated price is at the floor, a firm cannot compete on the anticipated price and can compete on the additional price only if it unshrouds—but because consumers who learn of the additional prices may not buy the product, the firm may find the latter form of competition unattractive. If this is the case for all firms, a profitable deceptive equilibrium exists. Otherwise, additional prices are unshrouded with probability 1, and firms earn zero profits.

The above condition for a firm to find unshrouding unattractive has some potentially important implications for when profitable deception occurs. First, if the product is socially wasteful (its value is lower than its production cost), then a firm that unshrouds cannot go on to profitably sell its product, so no firm ever wants to unshroud. Perversely, therefore, in a socially wasteful industry 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

identified and reached. Investigating trade and portfolio data from a large German bank, for example, Hackethal *et al*. (2010) document that "bank revenues from security transactions amount to €2,560 per customer per year" (2.4 % of mean portfolio value), a figure likely well above the marginal cost of serving a customer. Similarly, based on a number of measures, including the 20 % average premium in interbank purchases of outstanding credit card balances, Ausubel (1991) argues that credit card companies make large profits. Ellison and Ellison (2009) describe a variety of obfuscation strategies online computer-parts retailers use, and document that such strategies can generate surprisingly large profits given the near homogeneity of products.

2. In our applications, we provide context-specific reasons for the price floor. In Heidhues *et al*. (2012a), we develop a general 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 (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), Ko (2012), and Grubb (2015) also analyse models with variants of our price-floor assumption.

capturing the entire market, so if there is such a firm only a non-deceptive, zero-profit equilibrium exists. Hence, because some firm earns low profits in an industry with many firms, entry into socially valuable industries makes these industries more transparent; and whenever deceptive practices survive in an industry with many firms, the industry is socially wasteful.

We next allow for sophisticated consumers—who understand but cannot avoid the additional price—as well as naive consumers to be present in the market. Consistent with the predictions of Gabaix and Laibson (2006) and Armstrong and Vickers (2012) that sophisticated consumers facilitate transparency and efficiency, we find that if the product is socially valuable and there are sufficiently many sophisticated consumers, only a non-deceptive equilibrium exists. In a socially wasteful industry, however, a deceptive equilibrium exists independently of the share of sophisticated consumers.

In Section 4, we consider a multi-product market with a superior and an inferior product. We show that under weak conditions, there is an equilibrium in which sophisticated and naive consumers self-separate into buying the former and the latter product, respectively, and—in contrast to received wisdom—sophisticated consumers exert no pressure to unshroud the inferior product's additional prices. Worse, because the superior product renders the inferior product socially wasteful in relative terms, it *guarantees* that profitable deception in the market for the inferior product can be maintained. This observation has a striking implication: bad products are sold and remain profitable in a seemingly competitive market not despite, but *exactly because* superior alternatives exist, and the addition of a superior *or* an inferior product can expand the scope for profitable deception. Furthermore, we establish that when there are many firms, there is no deceptive equilibrium in which sophisticated consumers buy the inferior product at a competitive price, and naive consumers buy the superior product at an above-competitive price. Hence, if we observe that a product is sold profitably and deceptively to consumers when a competitively priced product is available, then the former product must be inferior.

In Section 5, we turn to extensions and modifications of our framework. We show that if a firm has market power in the superior product market, then it may have an incentive to unshroud to attract naive consumers to itself, especially if it sells mostly the superior product. But if its market power is small and unshrouding is costly, then the extent to which it educates consumers is limited. We also establish that if consumers initially underappreciate the additional price because they are unaware of a valuable add-on they want to spend on, then only an unshrouded prices equilibrium exists. In contrast, overestimation of a deceptive product's value can only strengthen the logic behind profitable deception.

In Section 6, we apply our framework to two important markets that have been invoked as conducive to hidden charges, mutual funds and credit cards, and argue that our results naturally organize a number of stylized facts from these industries. In our model of mutual funds, firms choose front loads investors understand and management fees investors ignore unless explained to them by a firm. A key section of the Investment Company Act of 1940, which prohibits the dilution of existing investors' shares in favour of new investors, prevents funds from offering discounts to new consumers, so that the front load must be nonnegative. Our model says that if active funds are inferior products, then shrouded active funds with high management fees sold to naive investors can—despite the possibility of unshrouding—coexist with low-fee passive funds sold to sophisticated investors, although there may be limited attempts at education by companies (*e.g.* Vanguard) specializing in passive funds. Going further, the observation that active funds are sold in a profitable deceptive way implies that they are indeed inferior—otherwise, an active fund would have an incentive to lower management fees somewhat and explain their importance to consumers, something we do not observe. Making things worse, because active funds are more profitable, firms are willing to spend resources to push these inferior products on consumers through advertising or commissions, and there is massive entry into the industry. But because competition among passive funds is fierce, only the most efficient passive funds can survive, so that passive funds tend to be larger than active funds.

In our model of credit cards, issuers make offers consisting of an anticipated price and an interest rate to naive time-inconsistent borrowers, who derive a convenience benefit from using a credit card and underestimate how much interest they will pay. In case of indifference, each consumer chooses whether to get or charge a card by lexicographically applying an exogenously given preference over cards. Then, any consumer who prefers a firm's card will get the card if the firm charges an anticipated price of zero. Any additional consumers a negative anticipated price attracts, therefore, will not use the firm's card, and hence these consumers are unprofitable. As a result, firms act as if they were facing a price floor of zero. Our model, therefore, says that credit cards may be profitably sold to naive consumers even when alternatives such as low-fee debit cards would create more social value, once again attracting many firms into an inferior industry and creating incentives to push an inferior product. Furthermore, the observation that unshrouding has not taken place implies that credit cards are indeed socially wasteful for many consumers—and in particular that excessive credit card borrowing is not only unanticipated, but also unwanted by consumers.

Finally, we briefly and informally lay out several other applications, including mortgages with changing terms, bank accounts, printers, and hotels, emphasizing that in some of these markets there may not be a binding floor on the anticipated price.

In Section 7, we discuss some policy implications of our model. Since deception is likely to be more common and have more adverse welfare consequences when the price floor is binding, policymakers should concentrate their regulatory efforts on industries with a binding price floor—of which supranormal profits is a telltale sign. In Section 8, we discuss the behavioural economics and classical literatures most closely related to our article. While a growing theoretical literature investigates how firms exploit naive consumers by charging hidden or unexpected fees, our article goes beyond all of the existing work in making the central prediction that socially wasteful and inferior products tend to be sold more profitably than better products, and identifying a number of additional implications of this insight. We conclude in Section 9 by pointing out important further questions raised by our model.

2. BASIC SINGLE-PRODUCT MODEL

$N \geq 2$ firms compete for naive consumers who have an outside option with utility $\underline{u}$, value each firm's product at $v > 0$, and are looking to buy at most one item. Firms play a simultaneous-move game in which they set anticipated prices f_n and additional prices $a_n \in [0, \overline{a}]$, and decide whether to costlessly unshroud the additional prices. A consumer who buys a product must pay both prices associated with that product—she cannot avoid the additional price. If all firms shroud, consumers make purchase decisions as if the total price of product n was f_n. If at least one firm unshrouds, all firms' additional prices become known to all consumers, and consumers make purchase decisions based on the true total prices $f_n + a_n$.[3] If consumers weakly prefer buying and

3. Our assumption on unshrouding—that a single firm can educate all consumers about all additional prices at no cost—is in the context of most real settings unrealistically extreme, especially if incurring the additional price depends partly on the consumer's own behaviour. This extreme case is theoretically useful for demonstrating that education often does not happen even if it is very easy. We show in our working paper (Heidhues *et al.*, 2012b) that if unshrouding is somewhat costly or reaches only a fraction of consumers, a deceptive equilibrium is more likely to occur, but our qualitative results do not change.

are indifferent between a number of firms, each of these firms gets a positive market share, with this share being $s_n \in (0,1)$ if consumers are indifferent between all firms.

Firm n's cost of providing the product is $c_n > 0$. We let $c_{\min} = \min_n\{c_n\}$, and—to ensure that our industry is competitive in the corresponding classical Bertrand model—assume that $c_n = c_{\min}$ for at least two firms. In addition, we assume that $v + \overline{a} > c_n$ for all n; a firm with $v + \overline{a} < c_n$ cannot profitably sell its product, so we think of it as not participating in the market. And to avoid uninteresting qualifications, we suppose that $v - \underline{u} \neq c_{\min}$.

Deviating from much of the literature, we impose a floor on the anticipated price: $f_n \geq \underline{f}$. We assume that $\underline{f} \leq v - \underline{u}$, which means that consumers are willing to buy the product if they perceive the total price to be $\underline{f}$. We also assume that $\underline{f} \leq c_{\min}$, so that firms are not prevented from setting a zero-profit total price.

Our main interest is in studying the Nash-equilibrium outcomes of the above game played between firms. While we fully characterize equilibrium outcomes in our main propositions, in discussing our results we focus on conditions for and properties of deceptive equilibria—equilibria in which all firms shroud additional prices with probability 1. Because no firm has an incentive to shroud if at least one firm unshrouds, there is always an unshrouded-prices equilibrium (an equilibrium in which unshrouding occurs with probability 1). When a deceptive equilibrium exists, however, it is more plausible than the unshrouded prices equilibrium for a number of reasons. Most importantly, we show in Appendix A that whenever a positive profit deceptive equilibrium exists in our model with no unshrouding cost, it is the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is. In addition, a positive profit deceptive equilibrium is preferred by all firms to an unshrouded prices equilibrium. Finally, for the lowest-priced firms, the strategy they play in an unshrouded prices equilibrium is weakly dominated by the strategy they play in a positive profit deceptive equilibrium.

To simplify our statements, we define the *Bertrand outcome* as the outcome of a Bertrand price-competition game with rational consumers and transparent prices: consumers buy if and only if $v - \underline{u} > c_{\min}$ and pay a total price of $c_{\min}$ whenever they buy, and firms earn zero profits.

3. PROFITABLE DECEPTION

This section analyses our basic model, and discusses key economic implications.

3.1. *Non-binding price floor*

First, we characterize equilibrium outcomes when the floor on the anticipated price is not binding. In our setting, this means that (even charging the maximum additional price) a firm cannot make profits if it chooses an anticipated price equal to the floor.

Proposition 1. (Equilibrium with non-binding price floor) *Suppose $\underline{f} \leq c_{min} - \overline{a}$. For any $\psi \in [0,1]$, there exist equilibria in which unshrouding occurs with probability ψ. If shrouding occurs in an equilibrium, then firms earn zero profits, and consumers buy the product from a most efficient firm, pay $f = c_{min} - \overline{a}, a = \overline{a}$, and get utility $v - c_{min}$. If unshrouding occurs in an equilibrium, then the Bertrand outcome obtains.*

Proposition 1 implies that if the price floor is not binding, there is always a deceptive equilibrium. Since in a deceptive equilibrium consumers do not take into account additional prices when choosing a product, firms set the highest possible additional price, making existing consumers valuable. Similarly to the logic of Lal and Matutes's (1994) loss-leader model as well as that of many switching cost (Farrell and Klemperer, 2007) and behavioural economics theories, firms compete aggressively for these valuable consumers *ex ante*, and bid down the anticipated price until they eliminate net profits. In addition, since with these prices a firm cannot profitably undercut competitors, no firm has an incentive to unshroud.

By Proposition 1, any equilibrium outcome is identical either to that in the above deceptive equilibrium or to that in an unshrouded prices equilibrium. If unshrouding occurs, consumers make purchase decisions based on the total price, so firms effectively play a Bertrand price-competition game, leading to a Bertrand outcome.

3.2. *Binding price floor*

We turn to analysing our model when the price floor is binding, assuming that $\underline{f} > c_n - \overline{a}$ for all n. This condition means that setting an anticipated price equal to the floor does not preclude a firm from making profits on its consumers.

We first identify sufficient conditions for a deceptive equilibrium to exist. As above, in such an equilibrium all firms set the maximum additional price $\overline{a}$. Then, since firms are making positive profits and hence have an incentive to attract consumers, they bid down the anticipated price to $\underline{f}$. With consumers being indifferent between firms, firm n gets market share s_n and therefore earns a profit of $s_n(\underline{f}+\overline{a}-c_n)$. For this to be an equilibrium, no firm should want to unshroud additional prices. If unshrouding occurs, consumers are willing to pay at most $v-\underline{u}$ for the product, so firm n can make profits of at most $v-\underline{u}-c_n$ by unshrouding. Hence, unshrouding is unprofitable for firm n if the following "Shrouding Condition" holds:

$$s_n(\underline{f}+\overline{a}-c_n) \geq v-\underline{u}-c_n. \tag{SC}$$

A deceptive equilibrium exists if condition (SC) holds for all n. Furthermore, since $s_n < 1$, condition (SC) implies that $\underline{f}+\overline{a} > v-\underline{u}$, so that in a deceptive equilibrium consumers are worse off than with their outside option. Proposition 2 summarizes this result, and characterizes equilibria fully[4]:

Proposition 2. (Equilibrium with binding price floor) *Suppose $\underline{f} > c_n - \overline{a}$ for all n.*

- *(I) If condition* (SC) *holds with a strict inequality for all n, then there is a deceptive equilibrium in which all firms offer the contract $(f_n, a_n) = (\underline{f}, \overline{a})$ with probability 1, consumers receive utility below their outside option, and firms earn positive profits. In any other equilibrium, unshrouding occurs with probability 1 and the Bertrand outcome obtains.*
- *(II) If condition* (SC) *is violated for some n, then unshrouding occurs with probability 1 and the Bertrand outcome obtains.*

The intuition for why firms might earn positive profits despite facing Bertrand-type price competition is in two parts. First, as in previous models and as in our model with a non-binding price floor, firms make positive profits from the additional price, and to obtain these *ex post* profits each firm wants to compete for consumers by offering a lower anticipated price. But once firms hit the price floor, they exhaust this form of competition without dissipating all *ex post* profits.

4. We ignore the knife-edge case in which no firm violates condition (SC) but condition (SC) holds with equality for some firm. The equilibrium outcomes characterized in Case I of Proposition 2 remain equilibrium outcomes in this case, but there can be additional equilibrium outcomes.

Secondly, since a firm cannot compete further on the anticipated price, there is pressure for it to compete on the additional price—but this requires unshrouding and is, therefore, an imperfect substitute for competition in the anticipated price. If a firm unshrouds and cuts its additional price by a little bit, a consumer learns not only that the firm's product is the cheapest, but also that all products are more expensive than she thought. If the consumer receives negative surplus from buying at the current total price (*i.e.*, if $\underline{f}+\overline{a}>v-\underline{u}$), this surprise leads her not to buy, so the firm can attract her by unshrouding only if it cuts the additional price by a discrete margin. Since this may be unprofitable, the firm may prefer not to unshroud.

As the flip side of the above logic, if purchasing at the total market price is optimal (*i.e.*, if $\underline{f}+\overline{a}\leq v-\underline{u}$), then despite the negative surprise a consumer is willing to buy from a firm that unshrouds and undercuts competitors' additional prices by a little bit, so in this case a deceptive equilibrium does not exist. Our model, therefore, says that deception—and positive profits for firms—must be associated with suboptimal consumer purchase decisions.[5]

Part II of Proposition 2 establishes that condition (SC) is not only sufficient, but also necessary for profitable deception to occur. To understand the rough logic, assume towards a contradiction that all firms shroud with positive probability. Then, a firm can ensure positive profits by shrouding and choosing prices $(\underline{f},\overline{a})$, so in equilibrium each firm earns positive expected profits. Now the key step in the proof shows an intuitively plausible fact: that when setting a high price in its equilibrium distribution of prices, a firm prefers to shroud, and can only hope to make positive profits if other firms also shroud. Conditioning on the event that all firms set such high total prices, incentives are as if firms were shrouding with probability 1. Exactly as above, therefore, firm n sets $(\underline{f},\overline{a})$ and earns $s_n(\underline{f}+\overline{a}-c_n)$. But firm n can earn $v-\underline{u}-c_n$ by unshrouding, so a firm that violates condition (SC) prefers to unshroud.

3.3. *Key economic implications*

We distinguish two cases according to whether the product is socially valuable or wasteful.

Non-vanishingly socially valuable product (there is an $\epsilon>0$ such that $v-\underline{u}>c_n+\epsilon$ for all n). In this case, the right-hand side of condition (SC) is positive and bounded away from zero. Hence, a firm with a sufficiently low s_n violates condition (SC)—since it earns low profits from deception, it prefers to attract consumers through unshrouding—and thereby guarantees that only a zero-profit unshrouded prices equilibrium exists. This implies that a deceptive equilibrium can exist if the number of firms is sufficiently small, but not if the number of firms is large—as some firm will then violate condition (SC). Furthermore, an increase in the number of firms can induce a regime shift from a high-price, deceptive equilibrium to a low-price, transparent equilibrium.[6]

5. Part I of Proposition 2 also states that whenever a deceptive equilibrium exists, there is also an equilibrium in which unshrouding occurs with probability 1 and—because prices are then transparent and the standard Bertrand logic holds—the Bertrand outcome obtains. As in the case of a non-binding price floor (Proposition 1), since unshrouding is costless, if a firm unshrouds it is (weakly) optimal for other firms to unshroud as well. Unlike in the case of a non-binding price floor, however, there is no equilibrium in which unshrouding occurs with an interior probability. Our proof of this claim elaborates on the following contradiction argument. Suppose that shrouding occurs with an interior probability, and consider a firm that unshrouds with positive probability and sets the highest price of any firm conditional on unshrouding. When setting this highest price, the firm earns profits only if all other firms shroud, and even then it earns at most $v-\underline{u}-c_n$. If it instead shrouds and sets $(\underline{f},\overline{a})$, then it earns at least $s_n(\underline{f}+\overline{a}_n-c_n)$ if all other firms also shroud. By condition (SC), therefore, the firm prefers to shroud with probability 1, a contradiction.

6. The reason for stating our result for *non-vanishingly* socially valuable products is that if the social value of a firm's product $(v-\underline{u}-c_n)$ could be arbitrarily small, then even a firm with low profits from deception might not be willing to unshroud. A precise condition for when unshrouding must occur is $N>(\underline{f}+\overline{a})/\epsilon$. Then, $s_n<\epsilon/(\underline{f}+\overline{a})$ for some n, and for this n we have $s_n(\underline{f}+\overline{a}-c_n)<\epsilon<v-\underline{u}-c_n$, in violation of condition (SC).

Socially wasteful product ($v-\underline{u}<c_n$ for all n). In this case, the right-hand side of condition (SC) is negative, while the left-hand side is positive. Hence, a deceptive equilibrium exists regardless of the industry's concentration and other parameter values:

Corollary 1. (Wasteful products) *Suppose $\underline{f}>c_n-\overline{a}$ and $v-\underline{u}<c_n$ for all n. Then, a profitable deceptive equilibrium exists.*

This perverse result has a simple logic: since a socially wasteful product cannot be profitably sold once consumers understand its total price, a firm can never make profits from unshrouding.

The different logic of socially valuable and socially wasteful industries in our model yields two potentially important further points. First, our theory implies that if an industry experiences a lot of entry and does not come clean in its practices, it is likely to be a socially wasteful industry. Secondly, our theory suggests a general competition-impairing feature in valuable industries that is not present in wasteful industries: to guarantee a positive profit deceptive equilibrium, each firm wants to make sure competitors earn sufficient profits from shrouding. This feature is likely to have many implications beyond the current article (as, for instance, for innovation incentives in Heidhues *et al.,* 2016), and implies that wasteful industries may sometimes be more fiercely competitive than valuable ones.

3.4. *Sophisticated consumers*

We now incorporate into our model consumers who observe—but cannot avoid—additional prices. We assume that the proportion of these sophisticated consumers is $\kappa\in(0,1)$, and that the price floor is binding ($\underline{f}>c_n-\overline{a}$ for all n).

Notice that in a deceptive equilibrium, sophisticated consumers do not buy the product: if they did, they would buy from a firm with the lowest total price, and such a firm would prefer to either undercut equal-priced competitors and attract sophisticated consumers, or (if there are no equal-priced competitors) to unshroud and attract all consumers. Hence, the total price of any firm exceeds $v-\underline{u}$. This implies that if firm n unshrouds, it maximizes profits by setting a total price equal to $v-\underline{u}$, attracting all naive and sophisticated consumers. Combining these considerations, unshrouding is unprofitable if

$$(1-\kappa)s_n(\underline{f}+\overline{a}-c_n)\geq v-\underline{u}-c_n, \tag{1}$$

and a deceptive equilibrium exists if and only if condition (1) holds for all n.

Condition (1) implies that if the product is socially valuable, the condition for a deceptive equilibrium to exist is stricter if sophisticated consumers are present than if they are not. Intuitively, since sophisticated consumers observe the additional price, they create pressure to cut the additional price—and by implication also to unshroud. This result is consistent with the insights of Gabaix and Laibson (2006) and Armstrong and Vickers (2012) that if the proportion of sophisticated consumers is sufficiently high, transparency and efficiency obtains. But if the product is socially wasteful, then as before a profitable deceptive equilibrium always exists. The reason is simple: sophisticated consumers never buy a socially wasteful product in equilibrium, so their presence is irrelevant—firms just exploit naive consumers.

4. MULTI-PRODUCT MARKETS

In this section, we study markets in which deceptive products and alternatives coexist. First, we modify our model above by assuming that each firm has a transparent product in addition to the potentially deceptive product we have described. Because we think of the alternative product as an endogenous outside option, we set consumers' utility from not buying, $\underline{u}$, to zero. Consumers' value for the transparent product is $w > 0$, and firm n's cost of producing it is $c_n^w \geq 0$, with at least two firms having the lowest cost $c_{\min}^w = \min_n\{c_n^w\}$. We posit that product w is socially valuable ($w - c_{\min}^w > 0$). Consumers are interested in buying at most one product. Firms simultaneously set the anticipated and additional prices for product v, the single transparent price for product w, and decide whether to unshroud the additional prices of product v. If consumers weakly prefer buying and are indifferent between a number of firms in the market for product v or w, firms split the respective market in proportion to s_n or s_n^w, respectively; and if consumers are indifferent between products v and w, a given positive fraction of them chooses product w. We define a Bertrand outcome as a situation in which consumers buy a social-surplus-maximizing product at a total price equal to the product's lowest marginal cost and firms earn zero profits; this is the outcome that would obtain in classical Bertrand price competition.

We now turn to the analysis of our model.

Proposition 3. **(Profitability of inferior products)** *Suppose product v is inferior to product w ($v - c_{\min} \leq w - c_{\min}^w$), $\underline{f} > c_n - \overline{a}$ for all n, and $v - \underline{f} > w - c_{\min}^w$. For any proportion κ of sophisticated consumers and any shares s_n, s_n^w, a deceptive equilibrium exists. In any deceptive equilibrium, firms sell product w to sophisticated consumers and earn zero profits on it, and they sell product v to naive consumers and earn positive profits on it. In any other equilibrium, unshrouding occurs with probability 1 and the Bertrand outcome obtains.*

The intuition is in two parts. First, $v - \underline{f} > w - c_{\min}^w$ implies that because naive consumers ignore the additional price, they mistakenly find the inferior product v more attractive than the superior product w. As a result, the two types of consumers separate, so—quite in contrast to the message of Section 3.4—sophisticated consumers do not create an incentive to unshroud the inferior product's additional prices. Secondly, if a firm unshrouded, naive consumers would simply switch to the superior product, so the firm would not attract any consumer to its inferior product. Using the logic of our single-product model, the superior product guarantees a deceptive equilibrium with positive profits from the inferior product by rendering the inferior product socially wasteful in comparison.

The sorting condition $v - \underline{f} > w - c_{\min}^w$ is likely to hold in many applications. For instance, since the price floor is often around zero and costs are positive, $c_{\min}^w > \underline{f}$ tends to hold, immediately implying the condition if $v \geq w$. Even if $v < w$, the condition holds as long as product w cannot be produced very cheaply. Intuitively, since product w is transparent, naive consumers underestimate the full price of product v but not of product w, generating a preference for product v.

In combination, Propositions 2 and 3 imply that for socially valuable products, the existence of a transparent superior product can expand the scope for profitable deception and make naive consumers worse off. Proposition 2 implies that if there is no other product around but there are many firms (or at least one small firm), unshrouding and marginal cost pricing obtain. But Proposition 3 implies that once a superior transparent product is present, shrouding can happen with any market structure.

Note that Proposition 3 holds for *any* market shares s_n, s_n^w for the two products. In particular, this means that not even a "specialist" in product w—a firm that sells mostly the superior product—has an incentive to unshroud. Intuitively, competition reduces the margin on the superior product to zero, so whether or not it unshrouds a specialist makes no money from the superior product.

To identify another economically important implication of our framework, we show that with many firms, the deceptive equilibrium in Proposition 3 relies on product v being inferior:

Proposition 4. (Profitable shrouding implies inferiority) *Suppose product v is non-vanishingly superior to product w (there is an $\epsilon > 0$ such that $v - c_n > w - c^w_{\min} + \epsilon$ for all n) and $\underline{f} > c_n - \overline{a}$ for all n. If N is sufficiently large, then there does not exist a deceptive equilibrium in which firms sell product w to sophisticated consumers and earn zero profits on it, and they sell product v to naive consumers and earn positive profits on it.*

Proposition 4 says that not only must profitable inferior products be sold in a deceptive way—a relatively unsurprising conclusion—but in a market with many firms, profitable shrouded products must be inferior. Intuitively, if the profitable shrouded product was superior, then a firm making low total profits would have an incentive to unshroud and sell that superior product more cheaply to all consumers.

To conclude this section, we establish that the logic of the deceptive equilibrium we have identified above survives unchanged if firms can charge an additional price also for the alternative product. Accordingly, we modify the model by assuming that rather than choosing a single transparent price for product w, each firm chooses an anticipated price $f^w_n \geq \underline{f}$ as well as an additional price $a^w_n \in [0, \overline{a}]$, and if at least one firm unshrouds, all additional prices of products v and w become known to all naive consumers. We also impose that $\underline{f} \leq c^w_{\min}$, so that firms are not prevented from setting a zero-profit total price on product w. We leave all other assumptions unchanged. Then:

Proposition 5. (Profitability of inferior products, cont.) *Suppose $\underline{f} > c_n - \overline{a}$ for all n. If product v is inferior to product w ($v - c_{min} \leq w - c^w_{min}$) and $v > w$, then there exists a deceptive equilibrium in which firms sell product w to sophisticated consumers and earn zero profits on it, and they sell product v to naive consumers and earn positive profits on it. If product v is non-vanishingly superior to product w (there is an $\epsilon > 0$ such that $v - c_n > w - c^w_{min} + \epsilon$ for all n) and N is sufficiently large, then such an equilibrium does not exist.*

Intuitively, in the equilibrium of interest it is sophisticated consumers who buy product w, and (irrespective of firms' unshrouding decisions) these consumers evaluate products based on the total prices. As a result, whether or not firms can charge an additional price for product w leaves our insights essentially unaffected. The only difference is that the sorting condition that guarantees a deceptive equilibrium now becomes $v > w$, ensuring that if the anticipated price of both products is at the floor, naive consumers prefer product v.

Proposition 5 implies that not only can the addition of a superior transparent product expand the scope for profitable deception, but—perhaps even more perversely—so can the addition of an *inferior* product. Furthermore, this expansion of products can create firms' profit base in (or shift it to) the new, inferior product.

Going beyond the setting of our model, the insights above have an immediate implication for the marketing of superior and inferior products: because the inferior product is more profitable, firms have an incentive to push it on consumers who may not otherwise buy it, further decreasing social welfare by expending resources to sell an inferior good. First, as formally analysed in Murooka (2013), firms may pay intermediaries to convince consumers to buy the inferior product. Secondly, firms may engage in persuasive advertising to induce demand for the inferior product, with the—to the best of our knowledge—novel implication that persuasive advertising is directed exclusively towards an inferior product. Thirdly, firms may actively inform consumers of the existence of the inferior product, yet not do the same for the superior product. Fourthly, firms may make costly (real or perceived) improvements to the inferior product to make it more attractive to consumers. In some markets, firms may burn all of their gross profits on pushing inferior products.

5. EXTENSIONS AND MODIFICATIONS

In this section, we discuss various extensions and modifications of our theory that are relevant in applications. Unless otherwise stated, we continue to assume that the price floor is binding.

5.1. *Market power*

We consider the effect of market power of a simple kind—we assume that firm n is strictly more efficient than competitors ($c_n < c_{n'}$ for all $n' \neq n$). In this case, the condition for when a deceptive equilibrium exists in our single-product model remains unchanged. Intuitively, competitors' costs do not affect their prices when shrouding occurs, so these costs do not play a role in determining whether firm n wants to unshroud.

Market power does have an interesting effect in our multi-product model. In particular, suppose that product w is transparent and superior to product v, and that firm n has market power in the market for product w: $c^w_{\min} - c^w_n \equiv M > 0$, where $c^w_{\min} \equiv \min_{n' \neq n} c^w_{n'}$. Unlike in Section 4, we do not impose a tie-breaking rule for product w. Then:

Proposition ***6 (Market power in the superior product)*** *Suppose $\underline{f} > c_{n'} - \overline{a}$ for all n', $v - c_{min} < w - c^w_{min}$, and $v - \underline{f} > w - c^w_n$. Then, there is a deceptive equilibrium in which all firms sell product v to naive consumers at $(\underline{f}, \overline{a})$ and firm n sells product w at price c^w_{min} to sophisticated consumers, if and only if $M \leq s_n(\underline{f} + \overline{a} - c_n)$.*

Proposition 6 implies that if firm n's market power is larger than its share of deceptive profits, then the deceptive equilibrium from our competitive model does not exist: firm n would rather unshroud and induce everyone to buy the superior product. Such a situation could arise, for instance, if firm n sells almost only the superior product ($s_n \approx 0$). But if firm n's market power is lower than its share of deceptive profits, then it prefers not to educate consumers. Similarly to Section 4, the intuition is that deception creates a large profit margin on the inferior product despite competition, and unshrouding the additional prices of this product only leads consumers to buy the less profitable superior one.

Going further, the extent of unshrouding can be even more limited if unshrouding is costly. Suppose firm n is a specialist in the superior product, but that market is quite competitive (M is small). Then, firm n is unwilling to educate naive consumers unless the unshrouding cost is small. And if it is relatively cheap to educate a low fraction of consumers, but much more expensive to educate a significant fraction, firm n would choose the former, limited education.

While Proposition 6 is stated assuming that product w is transparent, its proof and logic are essentially unchanged if firms can charge a shrouded additional price also for product w. As in Section 4, the only difference is that the sorting condition becomes $v > w$ instead of $v - \underline{f} > w - c^w_n$, so that firms cannot profitably sell product w to naive consumers without unshrouding.

Market power also has an interesting effect in our single-product model with a non-binding price floor when firms are asymmetric in their ability to exploit consumers, and the product is socially valuable. Suppose $c_n < \min_{n' \neq n} c_{n'}$, and denote by $\overline{a}_{n'}$ the maximum additional price firm n' can charge. Then, in a deceptive equilibrium the lowest $f_{n'}$ firm n' is willing to charge equals $c_{n'} - \overline{a}_{n'}$, so firm n cannot make profits greater than $\min_{n' \neq n}(c_{n'} - \overline{a}_{n'}) - (c_n - \overline{a}_n)$. If it unshrouds, however, firm n can make $\min_{n' \neq n} c_{n'} - c_n$, so if the latter is greater a deceptive equilibrium does not exist. Intuitively, if one firm has lower marginal cost but another can impose a higher additional price, then the former firm has an incentive to unshroud to gain a competitive edge. This suggests that even when the price floor is not binding, a deceptive equilibrium is more likely to occur with socially wasteful products.

5.2. *Misprediction of value*

We now consider the possibility that naive consumers misperceive a product's value in addition to its price. We begin with our single-product model. Paralleling our setup for additional prices, we assume that if all firms shroud, consumers perceive the value of the product to be $\tilde{v}$, and if at least one firm unshrouds, consumers correctly understand the value to be v. We suppose that $\max\{v, \tilde{v}+\overline{a}\} > c_n$ for all n; a firm with $\max\{v, \tilde{v}+\overline{a}\} < c_n$ cannot profitably sell its product, so we can think of it as not participating in the market. We also impose that $\underline{f} \leq \min\{\tilde{v}, v\} - \underline{u}$; this ensures that whether or not unshrouding occurs, consumers are willing to buy the product if they perceive the total price to be $\underline{f}$.

Under the above assumptions, the proof of Proposition 2 applies unaltered. This observation allows us to discuss some implications of consumer misprediction of value. An economically relevant possibility is one where consumers *underestimate* the value ($\tilde{v} < v$).

Corollary 2. (Unanticipated valuable add-on) *If $v-\tilde{v} \geq \overline{a}$, then in any equilibrium unshrouding occurs with probability 1.*

The condition $v-\tilde{v} \geq \overline{a}$ holds if consumers are completely unaware of a valuable add-on that they can purchase after getting the product. In this case, a deceptive equilibrium does not exist. Intuitively, if consumers are unaware of something they would be happy to pay for, then unshrouding cannot make them less willing to buy the product, so unshrouding and undercutting competitors is a profitable deviation from a candidate deceptive equilibrium.

Now suppose that consumers overestimate the product's value ($\tilde{v} > v$). Holding other parameters constant, this does not affect condition (SC), and hence does not affect whether a deceptive equilibrium exists in our model.[7]

Finally, we discuss value misperceptions in our multi-product model. The results of Proposition 3 are only strengthened if consumers' perceived value for the inferior product when shrouding occurs is $\tilde{v} > v$. Then, the sorting condition becomes weaker—$\tilde{v}-\underline{f} > w-c^w_{\min}$—as naive consumers are more prone to choosing the inferior product. If this weaker sorting condition holds, a deceptive equilibrium exists for the same reason as in Proposition 3: given that a superior product is available at marginal cost, no firm can make money once consumers perfectly understand the nature of the products, so no firm has an incentive to unshroud.

5.3. *Downward-sloping product demand*

We analyse a variant of our single-product model of Section 2 in which there is heterogeneity in v. We normalize $\underline{u}$ to 0 and denote the demand curve induced by the distribution of naive consumers' valuations by $D(f)$. We suppose that $D(\underline{f}) > 0$ and that there is a choking price (*i.e.*, an f' such that $D(f)=0$ for $f \geq f'$). Then, as an analogue of Proposition 2, a deceptive equilibrium with prices $(\underline{f}, \overline{a})$ exists if and only if

$$s_n D(\underline{f})(\underline{f}+\overline{a}-c_n) \geq \max_{f\in[\underline{f},\underline{f}+\overline{a}]} D(f)(f-c_n) \text{ for all } n. \tag{2}$$

7. Going slightly beyond our assumptions, however, consumers' overestimation of value can facilitate the creation of a socially wasteful industry by making consumers more willing to buy the product. In particular, if $v-\underline{f} < \underline{u} < \tilde{v}-\underline{f}$, then a profitable deceptive equilibrium exists when consumers perceive the value to be $\tilde{v}$, but not when they perceive the value to be v.

The left-hand side of inequality (2) is the profit firm n earns in the candidate deceptive equilibrium, while the right-hand side is the profit the firm earns when unshrouding and optimally undercutting competitors. To understand why unshrouding can be unattractive, note that unshrouding and slightly undercutting competitors leads consumers with values between $\underline{f}$ and $\underline{f}+\overline{a}$ not to buy, discretely reducing industry demand. This means that the optimal unshrouding strategy discretely reduces the markup (from $\underline{f}+\overline{a}-c_n$ to $f-c_n$) or discretely reduces demand (from $D(\underline{f})$ to $D(f)$), or both, undermining the standard Bertrand argument. Furthermore, if the product is socially wasteful to produce ($D(c_{\min})=0$), the right-hand side of inequality (2) is zero, so a deceptive equilibrium always exists. But if the product is socially valuable and there are sufficiently many firms, at least one firm would choose to unshroud, eliminating the deceptive equilibrium.[8]

6. APPLICATIONS

6.1. *Mutual funds*

We use our model to organize several observations about the mutual fund market, assuming that naive investors understand front loads but do not understand management fees.[9] To justify our bounds on the anticipated and additional prices, we begin by briefly discussing some institutional features. The Investment Company Act of 1940 in the U.S. regulates the trading prices and accounting of mutual funds. A fund must calculate its net asset value (NAV) based on market valuations, and including investment advisory fees and other expenses to date, at least once daily. Management fees are subject to fiduciary duty, which we interpret as putting a legal bound on the management fee.[10] The price of a share is simply the NAV divided by the number of shares, and investors can buy or redeem shares only at the per-share NAV.

8. Consider also what happens when there are sophisticated consumers in the population who are not separated by a superior transparent product, and who are heterogeneous in v. So long as a positive fraction of sophisticated consumers buys the product despite their knowing about the high additional price, a cut in the additional price attracts all these sophisticated consumers, so that an arbitrarily small fraction of these consumers induces some competition in the additional price. Whenever shrouding can be maintained, however, firms' profits are *not* driven to zero because—similarly to the "captive" consumers in Shilony (1977) and Varian (1980)—naive consumers provide a profit base that puts a lower bound on firms' total profit. Furthermore, it is clear that these profits can be sufficient to deter unshrouding.

9. Different types of empirical evidence suggest that—despite disclosure regulations—investors do not fully understand management fees, and that they appreciate front loads better. Using a natural experiment in India, Anagol and Kim (2012) show that mutual fund investors are less sensitive to amortized "initial issue expenses" than to otherwise identical "entry loads" paid up front. In a laboratory experiment on portfolio choice between S&P 500 index funds with different fees, Choi *et al*. (2010) find that subjects, including Wharton MBA students, Harvard undergraduates, and Harvard staff members, overwhelmingly fail to minimize fees. Using administrative data from the privatized Mexican Social Security system, Duarte and Hastings (2012) show that investors were not sensitive to fees in choosing between funds, leading to high management fees despite apparently strong competition. Barber *et al*. (2005) find a strong negative relationship between mutual fund flows and front loads, but no relationship between flows and operating expenses, and Wilcox (2003) also presents evidence that the overwhelming majority of consumers overemphasize loads relative to expense ratios. Based on a survey of 2,000 mutual fund investors, Alexander *et al*. (1998) report that only 18.9% could provide any estimate of their largest mutual fund's expenses, although 43% claimed that they knew this information at the time of purchase.

10. See Section 36 of the Investment Company Act. The Act does not impose a specific maximum fee. Nevertheless, fiduciary duty means that a manager must act in a way that benefits investors, and a management fee that leads an investor into (say) an almost certain loss relative to a benchmark clearly violates this restriction. Indeed, some lawsuits have successfully challenged excessive management fees on the basis of violating fiduciary duty: http://www.pionline.com/article/20111212/PRINT/312129971/wal-mart-merrill-lynch-settle-401k-fee-suit (accessed 24 May 2016).

In addition, a fund can impose purchase fees (front loads).[11] The language of the rules interpreting the Investment Company Act, and in particular the phrase and definition of purchase fees, strongly suggests that these fees are intended to be nonnegative. Furthermore, the same non-negativity constraint follows directly from Section 22(a) of the Act, which requires funds to "eliminat[e] or reduc[e] ... any dilution of the value of other outstanding securities" in favour of new consumers. One simple interpretation of this "no-dilution" rule is that new investors cannot be treated better than existing investors. Since both pay the management fees, an existing investor would strictly prefer to be treated as a new investor if the front load was negative, so a negative front load violates no dilution. What gives us extra confidence that this interpretation is correct is that the regulations on pricing were adopted as a way to interpret the no-dilution section of the Act. Indeed, consistent with a hard price floor of zero, we are unaware of any mutual fund that offers cash back, discounts, or other financial incentives to attract new investors.

Given the above institutional features, our multi-product model can be mapped directly to the mutual fund market in the following way. Suppose $N \geq 2$ firms offer active and passive funds to investors. Investors can invest in one active or one passive fund in period 0, with the investment paying off in period 1. Active and passive funds pay returns gross of fees of v and w, respectively. Funds simultaneously choose front loads and management fees (f_n, a_n) for their active funds and front loads and management fees (f_n^w, a_n^w) for their passive funds. The no-dilution rule imposes a floor $\underline{f} = 0$ on the front loads, and fiduciary duty imposes an upper bound $\overline{a}$ on the management fees.[12] The cost of serving an investor is c_n for an active fund and c_n^w for a passive fund. Because the cheapest passive funds are very cheap, for simplicity we assume that $c_{\min}^w = \min_n \{c_n^w\} = 0$. In addition, we assume that the highest permissible management fee can cover the cost of active funds $(\overline{a} > c_n)$.[13]

Crucially, we assume that active funds generate a higher gross value than passive funds $(v > w)$. There are at least two reasons that this may be the case. First, it might be that active funds simply produce higher gross returns. In fact, even if this is not the case in reality $(v \leq w)$, the analysis in Section 5.2 shows that the same results obtain if investors believe that it is $(\tilde{v} > w)$. Secondly, similarly to the notion of "money doctors" in Gennaioli *et al.* (2015), investors might get an emotional benefit from leaving their money with a specific person, adding an "emotional return premium" to actual returns.

Mapping our results to the mutual fund application allows us to reconcile a number of prominent features of the industry. To make our basic points, we take the perspective accepted by many researchers and observers—but debated by some others—that because few mutual fund managers can persistently and significantly outperform the market, most active funds are

11. See rules 2a-4 and 22c-1 adopted under the Investment Company Act, or the summaries by the Securities and Exchange Commission available at http://www.sec.gov/answers/nav.htm and http://www.sec.gov/divisions/investment/icvaluation.htm (accessed 24 May 2016). A fund can also impose redemption fees (rear loads), but since these are not important for our arguments, we do not incorporate them in our model.

12. An alternative reason for an upper bound on the management fees that can be collected is that higher management fees generate lower performance, and lower performance—independently of whether a consumer realizes the importance of fees—can trigger dissatisfaction, redemption, and finding alternative investment opportunities. This more elaborate variant of our model generates the same logic as that above.

13. Given that $\underline{f} = 0$, this is equivalent to the assumption that the price floor is binding. This assumption is consistent with the observation that many active funds have a front load equal to the price floor of zero. Specifically, in the U.S. market our model best fits level-load and no-load mutual fund shares ($6,748 billion in total net assets in 2012 according to the ICI fact book 2013, Figure 5.11), which generally have zero front loads. Front-end load shares ($1,881 billion in total net assets in 2012) have positive front loads, and are, therefore, a less obvious fit for our model with a binding price floor. But the typical such fund uses the front load largely to pay intermediaries' commissions, so the anticipated price collected by the fund is often at zero and hence again at the binding floor.

inferior to passive funds ($v-c_{\min} \leq w-c^{w}_{\min}$).[14] Because we also have $\underline{f} > c_n - \overline{a}$ for all n and $v > w$, the first part of Proposition 5 implies that an arbitrarily large number of profitable active funds can coexist—despite the possibility of investor education—with less profitable and for consumers higher-value passive funds. Intuitively, because active funds have higher gross value ($v > w$), naive investors prefer active funds to passive funds when both have zero front load and shrouding occurs. And because active funds are inferior on net ($v-c_{\min} \leq w-c^{w}_{\min}$) and passive funds are available to investors at cost, a firm cannot profitably sell anything if it unshrouds, so no firm wants to unshroud.

Assuming entry costs are sufficiently low, the fact that the market can support an arbitrary number of profitable active funds means that active funds should attract massive entry. Long-term trends in the industry bear out this prediction: since 1975, just before low-cost index funds were introduced, the number of active funds increased from 426 to 7,223.[15] In a classical model, inferior active funds should of course experience massive exit rather than entry—and indeed, for this reason the growth of active funds is often considered a puzzle (Gruber, 1996; French, 2008).

Beyond naturally explaining the co-existence of profitable active and less profitable passive funds under the assumption that active funds are inferior, our theory also contributes to the debate of whether they are inferior. Specifically, the second part of Proposition 5 says that if active funds were superior, then an equilibrium in which profitable active funds shroud does not exist—because a small active fund would be better off educating consumers about the importance of management fees and attracting a large pool of investors by slightly lowering these fees. The lack of such education efforts, therefore, provides additional support to the hypothesis that active funds are inferior.

While profitable active funds have no incentive to acquire consumers through education, they have strong incentives to acquire consumers through advertising and paying commissions to intermediaries. Exactly in line with this prediction, Mullainathan *et al*. (2011) document that firms tend to push expensive active funds through intermediaries. Consistent with the perspective that this hurts consumers, Bergstresser *et al*. (2009) find that broker-sold funds deliver lower risk-adjusted returns than do direct-sold funds.

In our basic model, not only do active funds not educate consumers, passive funds do not do so, either. As our results and discussion in Section 5.1 imply, however, if a firm mostly sells passive funds and has market power, it has at least some incentive to educate consumers. Consistent with this prediction, Vanguard has tried to educate consumers about the advantages of passive funds for decades.[16] At the same time, our model says that if a firm's market power is limited—which seems to be the case for Vanguard given its low margins—then its incentive to educate is also limited. Indeed, Vanguard prides itself on advertising much less than most competitors. For instance, it never charges 12b-1 fees—fees that can be used for marketing and distribution—and in 1999, when it had $442 billion under management, its advertising budget was a mere $8 million.[17]

14. See, for example, Gruber (1996), Carhart (1997), Kosowski *et al*. (2006), and French (2008). Berk and Green (2004) and Berk and van Binsbergen (2012), however, argue that the evidence is consistent with the hypothesis that most managers can outperform the market, and that active funds are not inferior to passive funds.

15. See ICI Fact Book 2013 (http://www.ici.org/pdf/2013_factbook.pdf), page 142. The number of funds in 2012 was 7,596, of which 373 were index funds. This number does not include mutual funds that own other mutual funds.

16. For instance, investors can gauge the implications of Vanguard's cost advantage at https://investor.vanguard.com/mutual-funds/low-cost (accessed 24 May 2016), at least if they find their way to this page.

17. Source: http://www.forbes.com/forbes/1999/0208/6303114a.html (accessed 24 May 2016). Kaufmann Fund, which was a hundredth of Vanguard's size, spent about the same amount on advertising. In addition, casual empiricism suggests that much of Vanguard's advertising is not about educating consumers regarding management fees, but about promoting the brand in non-informative ways (see, *e.g.* the "Vanguarding" campaign, http://www.nytimes.com/2010/03/15/business/media/15adco.html, accessed 24 May 2016).

Finally, our model predicts that—it being sold above cost—all firms offer an active fund, but—it being sold competitively—only the most efficient firms offer passive funds. To the extent that there is cost heterogeneity across firms, this implies that active funds tend to be smaller than passive funds. Although some alternative models may make the same prediction, this prediction is consistent with casual evidence that low-cost passive funds tend to be larger than active funds. For instance, while index funds hold 10% of mutual fund assets, seven of the ten largest mutual funds are index funds.[18]

6.2. *Credit cards*

We apply our framework to a credit card market in which borrowers have a naive taste for immediate gratification and hence underestimate how much interest they will pay.[19] We develop a market-specific model based on Heidhues and Kőszegi (2010), show that it simplifies to the reduced-form model of Section 2, and argue that our results help understand some features of the industry. We make a number of stark assumptions to simplify our model, but the insights are robust to many alternative specifications.

Issuers and consumers interact over four periods, $t=0,1,2,3$. In period $t=0,1,2$, a consumer aims to maximize total utility $u_t/\beta+\sum_{\tau=t+1}^{3}u_\tau$, where u_τ is instantaneous utility in period τ. The short-term discount factor β is distributed in the population according to the cumulative distribution function Q on $[0,1]$, where $\overline{a}\equiv\max_{\beta\in(0,1]}Q(\beta)(1/\beta-1)$ exists.[20] We posit that the consumer's period-0 total utility is relevant for welfare evaluations.

Issuer n's cost of managing a consumer's account is $c_n<\overline{a}$, and all issuers acquire funds at an interest rate of zero. In period 0, issuers choose anticipated prices (*i.e.*, salient charges such as annual fees) $f_n\in\mathbb{R}$ to be paid in period 3, and a gross interest rate $R_n\geq 1$. Simultaneously with its pricing decision, each firm decides whether to unshroud the costs associated with credit card use. If no firm unshrouds, consumers believe that they have $\beta=1$. If a firm unshrouds, consumers learn that their β's are drawn from $Q(\cdot)$, but they do not learn their individual β's.

Observing issuers' offers, in period 0 a consumer decides whether to get a credit card, valuing its future convenience use at v. The consumer's utility from her outside option, such as using cash or a low-fee debit card tied to her existing bank account, is $\underline{u}<v$.[21] A consumer can get multiple credit cards, but needs only one for convenience use, and gets other card(s) only if she sees a strict further benefit. If a consumer gets at least one credit card, she runs up charges of 1 in period 1. If a consumer is indifferent between multiple credit card deals, she lexicographically applies an exogenously given preference over cards to decide which one to get and charge. If she charges card n, she decides in period 2 whether to repay 1 to firm n immediately, or repay R_n in

18. For the top mutual funds, see http://www.marketwatch.com/tools/mutual-fund/top25largest (accessed 24 May 2016). The ICI Fact Book 2013 (http://www.ici.org/pdf/2013_factbook.pdf, accessed 24 May 2016) reports that index funds held \$1.3 trillion of the industry's \$13 trillion in net assets in 2012.

19. A substantial body of research in behavioural economics documents that individuals often have a taste for immediate gratification (Laibson *et al*., 2007; Augenblick *et al*., 2015, for instance), that such taste is associated with credit card borrowing (Meier and Sprenger, 2010), and that individuals may be naive about their taste (Laibson *et al*., 2007; Skiba and Tobacman, 2008). And consistent with the specific hypothesis that consumers underestimate costly borrowing on their credit cards, Ausubel (1991) finds that consumers are more responsive to the teaser rate in credit card solicitations than to the post-introductory interest rate, even though the latter is more important in determining the total amount of interest they will pay.

20. For presentational simplicity, we use a specification in which β divides current utility instead of the (behaviourally equivalent) more conventional specification in which β multiplies future utility. A sufficient condition for $\max_{\beta\in(0,1]}Q(\beta)(1/\beta-1)$ to exist is that $Q(\cdot)$ is continuous and $\lim_{\beta\to 0}Q(\beta)/\beta=0$.

21. Even if the consumer already has a debit card, her gross value from transacting with a credit card may be higher due to associated perks such as the grace period, airline miles, or free rental car insurance.

period 3. A unit of credit card spending generates a unit of instantaneous utility, and repaying a unit of outstanding debt costs a unit of instantaneous utility.

We now show that the above model simplifies to our reduced-form model. To facilitate the statement, we say that outcomes in two Nash equilibria are the same if the two equilibria differ at most in prices at which consumers do not buy.

Lemma 1. *The sets of Nash-equilibrium outcomes when f_n is unrestricted ($f_n \in \mathbb{R}$) and when f_n is restricted to be nonnegative ($f_n \geq 0$) are the same. Furthermore, if f_n is restricted to be nonnegative, then defining $a_n \equiv Q(1/R_n)(R_n - 1)$, the game is strategically equivalent to (i.e., for any action profile generates the same payoffs as) our reduced-form model with $\underline{f} = 0$.*

The key observation is that to avoid unprofitable consumers, issuers act as if they were facing a floor of zero on the anticipated price. Intuitively, although setting $f_n < 0$ rather than $f_n = 0$ induces more consumers to get issuer n's card, it does not prevent them from getting and using an alternative card they prefer. Since issuer n's profits derive from consumers using rather than just getting its card, the consumers it attracts with the negative anticipated price are unprofitable.

Given that $f_n \geq 0$ for all n, each consumer gets at most one card. The resulting interest payment on the card, then, acts as an additional price in that it adds to the issuer's profits and lowers the consumer's utility by the same amount. Furthermore, if shrouding occurs, a consumer expects to repay early, so she expects to pay no additional price, whereas if unshrouding occurs, she fully understands how much she will pay in expectation. This logic also makes clear that beyond interest payments, a late fee or any other unanticipated fee that transfers utility from consumers to issuers can also serve as an additional price.[22]

Our model provides several pieces of insight regarding the credit card market. To begin our discussion, we suppose that credit cards generate lower social value than the outside option ($v - c_{\min} \leq \underline{u}$). Since the consumer derives no value from the credit feature of the credit card and a debit card also supplies the convenience feature, we would expect this to be the case if the consumer's outside option is using her existing debit card. Then, Proposition 2 says that a deceptive equilibrium exists. Furthermore, pricing in the deceptive equilibrium is consistent with a pattern documented by DellaVigna and Malmendier (2004): that most popular credit cards in the U.S. offer a \$0 annual fee and a high regular interest rate.[23] Although DellaVigna and Malmendier's (2004) theory already explains the combination of low anticipated prices and high additional prices, our model also explains why the anticipated prices are at or close to zero.

The deceptive equilibrium in Proposition 2 also naturally accounts for the profitability of credit card issuing documented by Ausubel (1991) and the continuous entry into the industry

22. Notice that in our framework, an issuer prefers to undercut competitors' nonnegative anticipated prices if it can require consumers to use only the issuer's card in exchange, thereby generating interest payments down the line. This pricing strategy is consistent with the observation that many credit cards offer usage-contingent perks such as cash back or airline miles. But while such a strategy avoids unprofitable naive consumers, it may disproportionately attract unprofitable sophisticated consumers, so that to avoid these consumers firms may again act as if they were facing a price floor. Formally, suppose that a fraction λ of consumers has $\beta = 1$ and hence does not pay interest, and for these consumers the card's convenience benefit is v' and the outside option has utility $\underline{u}'$, with $v' - \underline{u}' < v - \underline{u}$. If firm n lowers f_n below $v' - \underline{u}'$, it attracts all these sophisticated consumers. If $\lambda(v' - \underline{u}' - c_{\min}) + (1-\lambda)(v' - \underline{u}' + \bar{a} - c_{\min}) \leq 0$, therefore, such a price cut is unprofitable.

23. There are two main categories of exceptions to the \$0 annual fee. First, charge cards—which require repayment within 30 days—often have an annual fee. This is consistent with the prediction of our multi-product model below, where sophisticated consumers who anticipate paying costly interest prefer to use a charge card with a higher anticipated price to a credit card with a lower anticipated price. Secondly, some cards offered to consumers with bad credit have an annual fee. This presumably compensates issuers for credit risk, a consideration that is outside our model.

emphasized by Evans and Schmalensee (2005). This is puzzling from the perspective of existing models given that there are no obvious diseconomies of scale in issuing, that from a consumer's perspective issuers (if not credit card networks) are likely to be almost undifferentiated, and that there are thousands of issuing banks (Bar-Gill, 2004).

In addition, our theory can be used to make inferences about consumers' value for credit card services from firm behaviour. Proposition 2 implies that with a socially valuable product and a large number of firms, unshrouding must take place. Hence, the observation that unshrouding has not taken place in this crowded industry suggests that credit cards are for many consumers not valuable. In line with this perspective, Laibson *et al.* (2007) estimate that owning a credit card makes the average household whose head has a high school degree but not a college degree worse off by the equivalent of paying \$2,000 at age 20. More specifically regarding the borrowing feature of credit cards, suppose that—unlike in our model with naive time-inconsistent borrowers—unanticipated borrowing generates unanticipated value for a consumer, for instance, by facilitating an unexpected important expense.[24] Then, Corollary 2 implies that unshrouding must take place with any number of firms. Hence, the observation that unshrouding has not taken place implies that expensive credit card borrowing is not only unanticipated, but—consistent with a time-inconsistent setup—also initially unwanted by consumers.

The above application of our single-product model is appropriate for situations in which a consumer is deciding whether to get a credit card in addition to a debit card she already owns. An alternative possibility is that the consumer is deciding between a credit card and a debit card when she initially owns neither. In this case, our multi-product model of Section 4 is appropriate. We can think of a credit card as the inferior product, and of a low-fee debit card as the superior product, with the convenience benefit of the two being the same ($v=w$).[25] Given that we derived the additional price on a credit card from unanticipated interest payments, and debit cards do not offer a credit option, we assume that debit cards can only be sold transparently. Proposition 3 implies that a deceptive equilibrium in which naive consumers get a credit card exists if $v-\underline{f}>w-c^w_{\min}$, or (using that $\underline{f}=0$ and $v=w$) $c^w_{\min}>0$—which is likely hold since issuing and administering a debit card is surely costly. In the deceptive equilibrium, sophisticated time-inconsistent consumers, who are aware that they would use the credit feature of a credit card excessively, choose a low-fee debit card for making payments.[26]

6.3. *Other potential applications*

6.3.1. Mortgages. For mortgages with changing repayment terms, such as pay-option mortgages, interest-only mortgages, and mortgages with teaser rates, many consumers underestimate the payments they will have to make once an initial teaser period ends.[27] In this

24. Formally, we can modify our setup by assuming that the consumer is time consistent, values the convenience benefit of the card at $\tilde{v}$, and can in period 2 spend 1 on something she values at $1+v-\tilde{v}$. She does not anticipate this need *ex ante*, and if she decides to spend, she cannot repay her credit card balance. This implies that she is willing to pay at most $\overline{a}=v-\tilde{v}$ to delay repayment.

25. The same conclusions hold if the convenience benefit of a credit card is higher ($v>w$).

26. In our credit card application, it is plausible to assume that some sophisticated consumers do not pay costly interest either because they are time consistent, or because they are wealthy enough to repay all charges within the grace period. In the deceptive equilibrium, such sophisticated consumers prefer the inferior product because—as emphasized in previous work starting from Gabaix and Laibson (2006)—they are cross-subsidized by naive consumers. If there are sufficiently few of these consumers, however, the price floor continues to bind and, therefore, the inferior product remains profitable. Going further, as we explain in Footnote 22, the existence of these consumers can create a price floor and hence ensure the profitability of the inferior product in the first place.

27. For instance, the Option Adjustable-Rate Mortgage (ARM) allows borrowers to pay less than the interest for a period, leading to an increase in the amount owed and sharp increases in monthly payments when the mortgage resets

environment, we think of the consumer's repayment costs assuming terms do not change as the anticipated price and of the increase in payments as the additional price. The additional price is limited by consumers' ability to walk away from the mortgage when payments increase. As we show in an earlier version of our article, a floor on the anticipated price can result because consumers would find an overly low anticipated price suspicious, and—concluding that "there must be a catch"—would not buy.[28] Finally, there is reason to believe that such mortgages are inferior or even socially wasteful: while a sharply increasing mortgage-payment schedule makes sense for the few consumers who can look forward to drastic increases in income (as argued by Cocco, 2013), the vast majority in this market would have been better served by traditional mortgages or not borrowing at all (Bar-Gill, 2008). Our models help explain why these bad products survived on a relatively competitive market despite the possibility of better educating consumers, and why—as documented by Agarwal and Evanoff (2013)—brokers steered borrowers eligible for superior products into these mortgages. Furthermore, the observation that marketers were not keen to explain to borrowers the exact nature of these mortgages provides an additional argument in support of the hypothesis that these products were indeed socially wasteful or inferior.

6.3.2. Bank accounts, hotels, and printers.. Our models and results also apply to markets with add-ons. In these applications, consumers buy a base product—account maintenance, a Las Vegas room, or a printer—and can then buy an add-on—such as costly overdrafting, gambling entertainment, or cartridges—whose cost they do not appreciate at the moment of purchase.[29] We can think of the base product's price as the anticipated price and of the utility loss from unexpectedly high add-on payments as the additional price. The additional price is limited by consumers' *ex post* demand response to add-on prices. Banks likely face a floor on the anticipated price close to zero for the same reason as do credit cards, explaining why account maintenance fees of many bank accounts are exactly zero (Armstrong and Vickers, 2012), and why these products are profitable. For hotels and printers, a price floor may arise from suspicion or other sources, but we do not know whether it is binding.

7. POLICY IMPLICATIONS

The suboptimal nature of equilibrium in our models raises the question of whether a social planner can improve outcomes. We discuss here a few issues related to potential policy interventions, though a full exploration of this difficult question is outside the scope of our article.[30]

to an amortizing payment schedule. See "Interest-Only Mortgage Payments and Payment-Option ARMs—Are They for You?", information booklet prepared for consumers by the Board of Governors of the Federal Reserve System, available at http://www.federalreserve.gov/pubs/mortgage_interestonly/mortgage_interestonly.pdf. Based on existing evidence (Cruickshank, 2000; Bucks and Pence, 2008; Gerardi *et al*., 2009; Woodward and Hall, 2012; Gurun *et al*., forthcoming) it is plausible to assume that consumers have a limited understanding of mortgages in general, and of these mortgages in particular. And some features of these mortgages, such as the Option ARM's one- or three-*month* introductory interest rate, seem to serve only the purpose of deceiving borrowers about the product's cost.

28. The notion that the price floor was binding—and hence lenders could not compete away higher additional prices with lower anticipated prices—is consistent with the observation that these products were often more expensive than available alternatives (Agarwal and Evanoff, 2013; Gurun *et al*., forthcoming). But because we need to know consumers' beliefs to identify this price floor, it is difficult to tell from primitives where the price floor is.

29. For more detailed arguments that these products can be interpreted as deceptive, see Gabaix and Laibson (2006) and Armstrong and Vickers (2012).

30. See Congdon *et al*. (2011) and especially Agarwal *et al*. (2009) for further discussions of the advantages and disadvantages of different types of policies aimed at mitigating the effects of consumer mistakes.

As a simple observation, the possibility in our single-product model that the industry shifts from deceptive to transparent pricing as the number of firms grows identifies a potential consumer-protection benefit of competition policy. Nevertheless, because in our multi-product model a firm specializing in the superior product has more incentive to educate consumers if it has market power than if it does not, competition is not uniformly beneficial.

An alternative, and in the context of deception more direct, way of intervening in the market is to regulate hidden charges.[31] To see the effects of such regulations in our setting, consider a policy that lowers the maximum additional price firms can charge from $\overline{a}$ to $\overline{a}' < \overline{a}$ in the range where the price floor is binding. If condition (SC) still holds with $\overline{a}$ replaced by $\overline{a}'$, firms charge $(\underline{f}, \overline{a}')$ in the new situation, so that the decrease in the additional price benefits consumers one to one. And if the decrease in the additional price leads to some firm violating condition (SC), the market becomes transparent and prices drop further. This provides a counterexample to a central argument brought up against many consumer-protection regulations: that its costs to firms will be passed on to consumers. Consistent with the prediction that consumers may benefit from regulatory controls on one component of the price (arguably hidden charges) that appears fully fungible with another component, Bar-Gill and Bubb (2012) and Agarwal *et al*. (2015) document evidence suggesting that the Credit Card Accountability, Responsibility, and Disclosure (Credit CARD) Act—while succeeding in lowering regulated fees—did not lead to an increase in unregulated fees or a decrease in the availability of credit, so that it lowered the total cost of credit to consumers.[32]

Beyond the main theme we have developed above—that in equilibrium socially wasteful and inferior products might be sold profitably to consumers—a further important general message emerges from our results: that deception is likely to be more widespread and economically harmful in the case of a binding than in the case of a non-binding price floor. With a binding price floor, firms strictly prefer deception to no deception, so they may take steps, such as engaging in "exploitative innovation" (Heidhues *et al*., 2016), to ensure that a deceptive equilibrium exists. In addition, less efficient producers may enter the market and sell positive amounts, and firms may spend on pushing and hence expanding the market share of inferior products. Finally, if there is some market power and firms are asymmetric, for socially valuable products a deceptive equilibrium is more likely with a binding price floor. These observations have two policy implications. First, if policymakers need to decide where to concentrate their regulatory efforts, they should—for both distributional and efficiency reasons—prioritize industries where the price floor is binding.

31. For example, the Credit CARD Act of 2009 limits late-payment, over-the-limit, and other fees to be "reasonable and proportional to" the consumer's omission or violation, thereby preventing credit card companies from using these fees as sources of extraordinary *ex post* profits. Similarly, in July 2008 the Federal Reserve Board amended Regulation Z (implementation of the Truth in Lending Act) to severely restrict the use of prepayment penalties for high-interest-rate mortgages. Regulations that require firms to include all non-optional price components in the anticipated price—akin to recent regulations of European low-cost airlines—can also serve to decrease $\overline{a}$.

32. It is important to note, however, that the direct regulation of additional prices is unlikely to be fully effective as a general approach to combating deception. As Agarwal *et al*. (2009) discuss in their setting, formulating *ex ante* guidelines for which charges are acceptable seems extremely difficult, and individually researching and approving each new financial product is very costly. One potential response to this challenge is for researchers to develop portable empirical methods for detecting additional prices and consumer mistakes from commonly available data. As a simple example, the empirical approach of Chetty *et al*. (2009) takes advantage of the general observation that if a consumer reacts to changes in financially identical price components differently, then she is not paying sufficient (relative) attention to some components of the price. Another response to the same challenge is to develop policies that do not rely on the social planner knowing which price components consumers underappreciate. For example, Murooka (2013) shows that when financial intermediaries motivated by commissions are used by firms to sell deceptive products, capping commissions, or requiring commissions to be uniform across products, can improve welfare even if the regulator cannot identify additional prices.

One sign of such an industry is the presence of surprisingly large profits given the level of competition. Secondly, in as much as the price floor arises for regulatory reasons—as in the mutual fund market—policymakers can attempt to relax the floor with improved regulation.

8. RELATED THEORETICAL LITERATURE

In this section, we discuss theories most closely related to our article. Our article goes beyond all of the existing work in making (and exploring further implications of) the central prediction that socially wasteful and inferior products tend to be sold more profitably than better products, and that deceptive products often must be inferior products. Indeed, in previous theories deceptive products are socially valuable.

In Gabaix and Laibson's (2006) model, firms sell a base good with a transparent price and an add-on with a shroudable price, and consumers buying the base good can avoid the add-on by undertaking costly steps in advance. Gabaix and Laibson's main prediction is that unshrouding the add-on prices can be unattractive because it turns profitable naive consumers (who fail to avoid the expensive add-on) into unprofitable sophisticated consumers (who avoid the add-on). Although the precise trade-off determining a firm's decision of whether to unshroud is different, we start from a similar insight, and use it to explore new implications.

Clarifying and adapting Gabaix and Laibson's theory, Armstrong and Vickers (2012) investigate a model of contingent charges in financial services, and apply it the U.K. retail banking industry. Consistent with our perspective, they argue that the "free if in credit" model—whereby firms charge nothing for account maintenance, and rely on contingent charges, such as overdraft protection, for revenue—can be naturally explained by the presence of naive consumers.[33]

In research complementary to ours, Grubb (2015) considers services (such as cellphone calls or bank-account overdraft protection) for which consumers may not know their marginal usage fee, and asks whether requiring firms to disclose this information at the point of sale increases welfare. If consumers correctly anticipate their probability of running into high fees, such price-posting regulation can actually hurt because it interferes with efficient screening by firms. If consumers underestimate their probability of running into fees, in contrast, fees allow firms to extract more rent from consumers, and price posting prevents such exploitation.[34]

Our theory is also related to the classical industrial-organization literature on markets in which firms sell a primary as well as a complementary good, consumers are locked in after the primary-good purchase, and consumers initially do not observe a firm's price for the complementary good. The main goal of this literature is to analyse the effect of market power in the complementary-good market on prices and welfare. Closely related to our benchmark Proposition 1, Lal and Matutes (1994) show that in competing for consumers, sellers offset the high price for the complementary

33. Relatedly, Piccione and Spiegler (2012) characterize how firms' ability to change the comparability of prices through "frames" affects profits in Bertrand-type competition. If a firm can make products fully comparable no matter what the other firm does—which is akin to unshrouding in our model and that of Gabaix and Laibson (2006)—the usual zero-profit outcome obtains. Otherwise, profits are positive. Piccione and Spiegler highlight that increasing the comparability of products under any frame through policy intervention will often induce firms to change their frames, which can decrease comparability, increase profits, and decrease consumer welfare. Investigating different forms of government interventions, Ko (2012) and Kosfeld and Schüwer (2011) demonstrate that educating naive consumers in Gabaix and Laibson's (2006) framework can decrease welfare because formerly naive consumers may engage in inefficient substitution away from the add-on.

34. Our theory also builds on a growing literature in behavioural industrial organization that assumes consumers are not fully attentive, mispredict some aspects of products, or do not fully understand their own behaviour. See, for instance, DellaVigna and Malmendier (2004), Eliaz and Spiegler (2006), Spiegler (2006a; 2006b), Laibson and Yariv (2007), Grubb (2009), and Heidhues and Kőszegi (2010).

good with a low price for the primary good. Nevertheless, Shapiro (1995), Hall (1997), and Borenstein *et al*. (2000) show that social welfare may not be maximized due to deadweight losses from consumers' reactions to the overly low price of the primary good and the overly high price of the complementary good. Our article has a different objective—analysing the scope for and implications of profitable deception in a competitive market—and shows that deception can lead to additional welfare losses from the systematic sale and pushing of inferior products.

Relatedly, in assuming that consumers can be induced to pay high additional fees once they buy a product, our theory shares a basic premise with the large literature on switching costs. But even if firms cannot commit to *ex post* prices and there is a floor on *ex ante* prices—so that positive profits obtain in equilibrium—our model's main insights do not carry over to rational switching-cost models. Most importantly, we are unaware of any rational switching-cost model that predicts the systematic sale of inferior products in competitive markets, and under the natural assumption that consumers know or learn product attributes, a firm is indeed better off selling a superior product.[35]

9. CONCLUSION

In our model, we have imposed exogenously that firms can unshroud additional prices. An important agenda for future research is identifying the kinds of education that convince naive consumers. This is likely to require a more elaborate understanding of naivete, as well as to open up further questions regarding who is able to influence consumers in equilibrium.

We have also taken the opportunity to deceive consumers—that is, the shroudable additional price component—as exogenous. In most real-world markets, however, someone has to come up with ways to hide prices from consumers, so that the search for deception opportunities can be thought of as a form of innovation. In our companion paper (Heidhues *et al*., 2016), we study the incentives for such "exploitative innovation", and contrast them with the incentives for innovation that benefits consumers. Here too we find a perverse incentive: because learning ways to charge higher additional prices increases the profits from shrouding and thereby lowers the motive to unshroud, a firm may have a strong incentive to make exploitative innovations *and* have competitors copy them. The possibility of exploitative innovation also adds caution to policies that directly regulate the maximum additional price: such a policy can greatly increase firms' incentive to make new exploitative innovations, and hence may have a small net effect.

APPENDIX

A. COSTLY UNSHROUDING

Consider the same game as in Section 2, except that each firm n has to pay $\eta \geq 0$ to unshroud additional prices.

Proposition 7. (Unique equilibrium with costly unshrouding) *Fix all parameters other than η, and suppose that $\underline{f} > c_n - \overline{a}$ for all n and a deceptive equilibrium exists for $\eta = 0$. Then, for any $\eta > 0$ there exists a unique equilibrium, and in this equilibrium all firms shroud and offer $(\underline{f}, \overline{a})$ with probability one.*

35. This point is immediate if consumers know the products' values at the time of original purchase. As an illustration of the same point when consumers learn product values only after initial purchase, suppose that there are two products with values v_H and v_L and costs c_H and c_L, respectively, product H is strictly superior ($v_H - c_H > v_L - c_L$), and the switching cost is k. Consider a firm who could sell either product to a consumer at an *ex ante* price of p followed by an *ex post* price of its choice. If the firm sells the superior product, it can charge an *ex post* price of $c_H + k$ in a competitive market, so it makes profits of $p + c_H + k - c_H$. If the firm sells the inferior product, then (as the consumer realizes that she can switch to the superior product) it can charge an *ex post* price of $c_H + k - (v_H - v_L)$, so it makes profits of $p + c_H + k - (v_H - v_L) - c_L$. It is easy to check that the firm prefers to sell the superior product.

Proposition 7 says that if a deceptive equilibrium exists in our basic model, then it is the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is. This provides one justification for our presumption that firms play a deceptive equilibrium whenever it exists.

To see the logic of Proposition 7, suppose that some firm unshrouds with positive probability. Notice that since $\eta > 0$, in order to unshroud a firm must make positive gross profits afterwards. Hence, no firm unshrouds with probability 1—as this would lead to Bertrand-type competition and zero gross profits. Now for each firm that unshrouds with positive probability, take the supremum of the firm's total price conditional on the firm unshrouding, and consider the highest supremum. At this price, a firm cannot make positive profits if any other firm also unshrouds. Hence, conditional on all other firms shrouding at this price, the firm must make higher profits from unshrouding than from shrouding. But this is impossible: if the firm has an incentive to shroud in this situation with zero unshrouding cost—which is exactly the condition for a deceptive equilibrium to exist—then it strictly prefers to shroud with a positive unshrouding cost.

B. PROOFS

B.1. *Proof of Proposition 1*

First, we establish that for any $\psi \in [0,1]$ there exists an equilibrium with the properties stated in the proposition. Let every firm offer the contract $f_n = c_n - \bar{a}, a_n = \bar{a}$ and unshroud with probability γ, where $(1-\gamma)^N = 1-\psi$. In this equilibrium, every firm makes zero profits independent of whether shrouding occurs or not, and—again independent of whether shrouding occurs or not—no firm can attract customers at a total price above its marginal cost. In what follows, we establish that any equilibrium has the properties specified in the proposition.

Suppose unshrouding occurs with probability $\psi = 1$. Then we have Bertrand competition in the total price, and standard arguments imply that the Bertrand outcome obtains.

We are thus left to consider the case in which unshrouding occurs with probability $\psi < 1$. Note first that no firm makes sales at a total price $f_n + a_n < c_{\min}$ with positive probability in equilibrium because such a firm could profitably deviate by moving all probability mass from total prices strictly below $c_{\min}$ to a total price of $c_{\min}$.

Next, we prove by contradiction that no firm makes sales at a total price $f_n + a_n > c_{\min}$ with positive probability. If this is not the case, then any most efficient firm, *i.e.*, any firm with cost $c_n = c_{\min}$, can copy this firm's strategy and thereby earn positive profits. Hence, any most efficient firm makes positive profits in equilibrium. Denote by $\pi_n^* > 0$ a most efficient firm n's equilibrium profits. Let $\bar{f}$ be the supremum of the anticipated prices of most efficient firms conditional on shrouding. Since $\underline{f} + \bar{a} \leq c_{\min}$, in order for a most efficient firm to make profits it must charge $f > \underline{f}$, so that $\bar{f} > \underline{f}$.

Suppose, first, that some most efficient firm n sets the anticipated price $\bar{f}$ with positive probability. Denote by $\hat{a}$ the supremum of firm n's additional price conditional on shrouding and setting an anticipated price of $\bar{f}$. Then, it cannot be the case that firm n earns positive expected profits when it sets $\bar{f}$ and shrouding occurs: this would only be possible if another most efficient firm also set $\bar{f}$ with positive probability when shrouding, but then firm n would benefit from minimally undercutting the anticipated price. Hence, conditional on charging $(\bar{f}, a_n)$ and shrouding, firm n can earn positive expected profits only when unshrouding occurs and all rivals charge a total price equal to or above $\bar{f} + a_n$ with positive probability. Thus, with positive probability all rivals must charge total prices weakly above $\bar{f} + \hat{a}$. This implies that if firm n charges $\bar{f} + \hat{a}$ with positive probability, then all rivals must charge total prices strictly above $\bar{f} + \hat{a}$ with positive probability; if rivals charged total prices equal to $\bar{f} + \hat{a}$ with positive probability, then firm n would strictly increase profits by minimally lowering its anticipated price.

Let $\bar{t}_{n'}$ be the supremum of the total prices of firm n' conditional on unshrouding. Let $\bar{t} = \max\{\bar{t}_{n'} | c_{n'} = c_{\min}\}$. With slight abuse of notation, say firm n' is the most efficient firm that achieves $\bar{t}$ when unshrouding. First, suppose firm n' sets $\bar{t}$ when unshrouding with positive probability. Because no firm other than n' can set a total price $\bar{t}$ with positive probability (otherwise, firm n' would prefer to minimally lower its total price), firm n' makes zero profits if any other most efficient firm unshrouds, so firm n' must earn positive profits when all other most efficient firms shroud. Hence, all most efficient firms $n'' \neq n'$ must with positive probability set a total price $f_{n''} + a_{n''} > \bar{t}$ when shrouding. Consider the supremum of total prices $\tilde{f} + \tilde{a}$ of a most efficient firm when shrouding. A firm $\tilde{n}$ setting this total price with positive probability makes zero profits when a most efficient firm unshrouds, so it must earn positive profits conditional on all most efficient firms shrouding, and can do so only when shrouding occurs. Hence, $\tilde{f} < \bar{f}$ and $\tilde{f}$ is with positive probability the lowest anticipated price. But then the most efficient firm setting $\bar{f}$ can profitably deviate by moving the probability mass from $\bar{f}$ to $\tilde{f}$ in such a way that the total price distribution remains unchanged (which it can do because $\tilde{f} + \tilde{a} > \bar{f} + \hat{a}$), a contradiction. Hence, no firm sets $\tilde{f} + \tilde{a}$ with positive probability. Let firm $\tilde{n}$ be a most efficient firm that achieves the supremum $\tilde{f} + \tilde{a}$ when shrouding. Consider a sequence of optimal total prices $\tilde{f}_k + \tilde{a}_k$ set by firm $\tilde{n}$ when shrouding that converges to $\tilde{f} + \tilde{a}$. Then, firm $\tilde{n}$'s expected profits conditional on unshrouding occurring go to zero, and hence it must earn positive expected profits conditional on shrouding occurring. This implies that there exists a subsequence $\tilde{f}_k$ converging to some $\tilde{f} < \bar{f}$ with the property that $\tilde{f}$ is with positive probability the lowest anticipated price. But then again the most efficient firm setting $\bar{f}$ can profitably deviate as above. We conclude that no firm sets $\bar{t}$ with positive probability.

Next, suppose that no firm sets $\bar{t}$ with positive probability. Consider a sequence $t_{n'} \to \bar{t}$ of optimal total prices by firm n' when unshrouding; then conditional on a most efficient firm unshrouding, firm n''s expected profits go to zero as $t_{n'} \to \bar{t}$; hence, to earn its equilibrium profit level, firm n''s expected profits conditional on all most efficient firms shrouding must be strictly positive and bounded away from zero. Hence, all other most efficient firms $n'' \neq n'$ must with positive probability set a total price $f_{n''} + a_{n''} \geq \bar{t}$ when shrouding. By the same argument as in the preceding paragraph, we arrive at a contradiction. Therefore, we conclude that no firm sets $\bar{f}$ with positive probability when shrouding.

Suppose, thus, that no most efficient firm charges $\bar{f}$ with positive probability when shrouding. Let firm n be a firm that achieves this supremum. Note that firm n's expected profits conditional on shrouding occurring go to zero as $f_n \to \bar{f}$. Let $\hat{a} = \limsup_{\epsilon \searrow 0}\{a|(f,a)$ optimal and $f > \bar{f} - \epsilon\}$. Consider a sequence (f_n^l, a_n^l) of optimal prices that converges to $(\bar{f}, \hat{a})$. For sufficiently high l, firm n must make expected profits conditional on unshrouding occurring and all rivals charging a total price at or above $\bar{f} + \hat{a}$. From now on, we follow the same arguments as the two preceding paragraphs (for the case in which some firm sets $\bar{f}$ with positive probability when shrouding) to establish (in the above notation) that $\tilde{f} < \bar{f}$ and $\tilde{f}$ is with positive probability the lowest anticipated price. Then for sufficiently high l, firm n's profits when charging an optimal (f_n^l, a_n^l) and shrouding occurring go to zero; but firm n could lower the anticipated price from f_n^l to $\tilde{f}$ and still charge the same total price $f_n^l + a_n^l$, increasing its expected profits conditional on shrouding occurring and not affecting it otherwise, contradicting the optimality of (f_n^l, a_n^l). This contradiction establishes that $\bar{f} > \underline{f}$ is impossible.

We have thus established that no firm sells a contract at a total price other than $c_{\min}$. Next, observe that consumers must buy the product at a total price of $c_{\min}$ with probability 1 if $v - \underline{u} > c_{\min}$; otherwise, all firms charge a price strictly greater than $c_{\min}$ with positive probability, and then a most efficient firm can make positive profits, which we showed above is impossible. Note also that if $v - \underline{u} < c_{\min}$ and unshrouding occurs, then consumers do not buy the product. This implies that if unshrouding occurs, the Bertrand outcome obtains.

Finally, we show that conditional on shrouding occurring, consumers must pay $f = c_{\min} - \bar{a}, a = \bar{a}$ with probability 1, and, therefore, get utility $v - \underline{u} - c_{\min}$. Since no firm makes sales at total prices below $c_{\min}$ with positive probability, no firm sells at an anticipated price below $c_{\min} - \bar{a}$ with positive probability. Now suppose that consumers buy at an anticipated price above $c_{\min} - \bar{a}$ with positive probability conditional on shrouding occurring. Then, there exists $\epsilon > 0$ for which consumers buy at an anticipated price greater than $c_{\min} - \bar{a} + \epsilon$ conditional on shrouding occurring with positive probability, and a most efficient firm could profitably deviate through shrouding and offering the contract $f' = c_{\min} - \bar{a} + (\epsilon/2), a' = \bar{a}$, a contradiction. ‖

B.2. *Proof of Proposition 2*

(I) Note first that if some firm unshrouds with probability 1, all other firms are indifferent between shrouding and unshrouding. Thus, an equilibrium in which unshrouding occurs with probability 1 always exists. In any such equilibrium consumers observe and take the additional price into account, so that our game reduces to a standard Betrand game in which the consumers' willingness to pay is $v - \underline{u}$. Hence, standard arguments imply that the Bertrand outcome obtains.

Now consider deceptive equilibria, *i.e.*, equilibria in which shrouding occurs with probability 1. In case firm n has a positive probability of sales in equilibrium, it must set $a_n = \bar{a}$ as otherwise it could increase its profits conditional on a sale by increasing a_n without affecting the probability of selling. The same argument as in the text establishes that if condition (SC) holds for all n, then there is a deceptive equilibrium in which all firms set $(\underline{f}, \bar{a})$. We now provide a formal argument for why firms set $\underline{f}$ with probability 1 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 all firms set the additional price $\bar{a}$. Note that by setting $f_n = \underline{f}$, firm n can guarantee itself a profit of $s_n(\underline{f} + \bar{a} - c_n) > 0$. As a result, no firm will set $f_n > v - \underline{u}$, because then no consumer would buy from it. Take the supremum $\bar{f}$ of the union of the supports of firms' anticipated price distributions. We consider two cases. First, suppose that some firm sets $\bar{f}$ with positive probability. In this case, all firms have to set $\bar{f}$ with positive probability; otherwise, a firm setting $\bar{f}$ would have zero market share and hence zero profits with probability 1. Then, we must have $\bar{f} = \underline{f}$; otherwise, a firm could profitably deviate by moving the probability mass to a slightly lower price. Secondly, suppose that no firm sets $\bar{f}$ with positive probability. Let firm n's price distribution achieve the supremum $\bar{f}$. Then, as f_n approaches $\bar{f}$, firm n's expected market share and hence expected profit approaches zero—a contradiction.

We now establish that if the strict inequality

$$s_n(\underline{f} + \bar{a} - c_n) > v - \underline{u} - c_n \tag{3}$$

holds for all n, in any equilibrium shrouding occurs either with probability 1 or zero. Suppose, towards a contradiction, that all firms shroud with probability strictly between 0 and 1. If all competitors shroud with positive probability, firm n can guarantee itself positive profits by shrouding itself and offering the contract $(\underline{f}, \bar{a})$, which attracts consumers since $v - \underline{u} \geq \underline{f}$ and makes positive profits since $\underline{f} + \bar{a} > c_n$. To earn positive profits when unshrouding, firm n must set a total price $t_n \leq v - \underline{u}$. 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 or more firms did so, then some firm could increase profits by moving this probability mass to slightly below $\hat{t}$. First, consider the case in which one firm puts positive probability mass on $\hat{t}$, and let this firm be firm n. Firm n earns zero profits conditional on some other firm unshrouding. Conditional on all others shrouding it earns less than

$$\hat{t}-c_n \leq v-\underline{u}-c_n < s_n(\underline{f}+\overline{a}-c_n),$$

and hence would increase its profits by deviating from unshrouding and charging the total price $\hat{t}$ to shrouding and charging $(\underline{f},\overline{a})$. Secondly, consider the case in which no firm puts positive probability mass on $\hat{t}$. Let n be a firm that achieves this supremum. Consider a sequence $t_k \to \hat{t}$ of optimal total prices by firm n when unshrouding, and denote the corresponding expected profits conditional on some other firm unshrouding by ϵ_k. Then, $\epsilon_k \to 0$. Notice that conditional on all other firms shrouding, firm n earns at most $t_k - c_n \leq v-\underline{u}-c_n$, so that the (unconditional) expected profit of firm n is less than $v-\underline{u}-c_n+\epsilon_k$, which by condition (3) is strictly less than $s_n(\underline{f}+\overline{a}-c_n)$ for a sufficiently small $\epsilon_k>0$. Hence, firm n is strictly better off shrouding and charging $(\underline{f},\overline{a})$—a contradiction. Thus, if condition (3) holds, unshrouding occurs with probability 1 or zero in equilibrium.

(II) We establish by contradiction that if condition (SC) is violated for some firm, then in any equilibrium additional prices are unshrouded with probability 1. Note again that if unshrouding occurs with probability 1, we have Bertrand competition in the total price, and hence the Bertrand outcome obtains. The proof that unshrouding occurs with probability 1 in this case proceeds in three steps.

Step (i): 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 the contract $(\underline{f},\overline{a})$, which attracts consumers since $v-\underline{u}\geq \underline{f}$ and makes positive profits since $\underline{f}+\overline{a}>c_n$.

Step (ii): All firms choose the anticipated 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 other than 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 $\overline{a}$ with probability 1. Take the supremum of firms' anticipated prices $\bar{f}'$ conditional on the total price being strictly higher than $\hat{t}$. Because consumers do not buy the product if the anticipated price is greater than $v-\underline{u}$ and firms must earn positive profits by (i), $\bar{f}'\leq v$. Note that $\bar{f}'+\overline{a}>\hat{t}$, so $\overline{a}>\hat{t}-\bar{f}'$.

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

If only one firm n' sets $\bar{f}'$ 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 (i) and is the only firm that sets $\bar{f}'$ with positive probability conditional on the total price being strictly higher than $\hat{t}$, every firm n'' other than n' sets its anticipated price strictly higher than $\bar{f}'$ 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 anticipated fee $f_{n''}>\bar{f}'$, $a_{n''}\leq \hat{t}-f_{n''}$ with positive probability. Since $\overline{a}>\hat{t}-\bar{f}'$, firm n'' can increase its profits by decreasing all prices $f_{n''}>\bar{f}'$ to $\bar{f}'$ and by increasing its additional price holding the total price constant—a contradiction. Next, suppose $n'\neq n$. Then, firm n shrouds, sets $f_n>\bar{f}'$ with positive probability and charges an additional price $a_n\leq \hat{t}-f_n$ with probability 1 when charging these anticipated prices. For almost all of these anticipated prices, firm n must earn strictly positive profits when shrouding occurs; otherwise firm n could unshroud with probability 1 and generate positive (and hence higher) profits when all rivals shroud and charge a total price above $\hat{t}$. Thus, firm n' shrouds and sets $f_{n'}\geq f_n>\bar{f}'$, $a_{n'}\leq \hat{t}-f_{n'}$ with positive probability. Then, firm n' can increase its profits by decreasing all prices $f_{n'}>\bar{f}'$ to $\bar{f}'$ and increasing its additional price holding the total price constant—a contradiction.

If no firm sets $\bar{f}'$ with positive probability, there exists firm n' that for any $\epsilon>0$ sets anticipated prices in the interval $(\bar{f}'-\epsilon,\bar{f}')$ with positive probability. As $\epsilon\to 0$, the probability that firm n' charges the highest anticipated price and the total price is strictly higher than $\hat{t}$ goes to one. Also, by the definition of $\hat{t}$, firm n' shrouds when setting these anticipated prices. Therefore, the profits go to zero 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 have established that $\bar{f}'=\underline{f}$.

Because $\bar{f}'=\underline{f}$, each firm $n'\neq n$ sets an anticipated price of $\underline{f}$ with probability 1 conditional on its total price being strictly higher than $\hat{t}$. Hence, $\underline{f}+\overline{a}\geq\hat{t}$. We now show that whenever shrouding, any firm $n'\neq n$ does not set anticipated 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 1 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 anticipated price strictly above $\underline{f}$. Since n' sets anticipated prices strictly above $\underline{f}$ with positive probability when shrouding, there exists an anticipated price $g' > \underline{f}$ such that firm n' sets prices above g' with positive probability. There cannot be a competitor whose anticipated price when shrouding falls on the interval $[\underline{f}, g']$ 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. Note that by the first sentence of this paragraph, firm n' sets $\underline{f}$ with positive probability when shrouding. But then, firm n' can raise its anticipated price from $\underline{f}$ to g' and increase profits—a contradiction. Thus, any firm $n' \neq n$ sets the anticipated price $\underline{f}$ with probability 1 when shrouding.

Now suppose that firm n charges an anticipated 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 and setting an anticipated price strictly greater than $\underline{f}$. 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 1 when shrouding.

Secondly, 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' anticipated prices $\bar{f}'$ conditional on the total price being greater than or equal to $\hat{t}$. The remainder of the proof is the same as above.

Step (iii): Additional prices are unshrouded with probability 1. 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 $\underline{t} \leq v - \underline{u}$. Take a firm that achieves the infimum. By (i), this firm earns positive profits. For any $\epsilon > 0$, take total prices below $\underline{t} + \epsilon$ of the firm. By unshrouding and setting $\underline{t} - \epsilon$ instead, the firm decreases its profits by at most 2ϵ when one or more other firms unshroud, but (since by (ii) all firms set an anticipated price of $\underline{f}$ when shrouding) discretely increases its market share if all other firms shroud. Hence, for sufficiently small $\epsilon > 0$ this is a profitable deviation—a contradiction. Secondly, suppose $\underline{t} > v - \underline{u}$. Take firm n that violates condition (SC). By (ii), firm n charges the anticipated price $\underline{f}$ whenever it shrouds. Note that firm n's profits are zero when a rival unshrouds, and its profits are at most $s_n(\underline{f} + \overline{a} - c_n)$ 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 $v - \underline{u} - c_n > s_n(\underline{f} + \overline{a} - c_n)$. ‖

B.3. *Proof of Proposition 3*

First, if unshrouding occurs with probability 1, all consumers are sophisticated and standard (Bertrand-competition) arguments imply that they buy a total-surplus-maximizing product at marginal cost.

Now suppose, towards a contradiction, that shrouding occurs with positive probability strictly less than one. Our proof that this is impossible has four steps. In Step (i) we show that if shrouding occurs with positive probability, firms earn positive profits. Step (ii), which is contained in Lemma 2, establishes that every firm sets an anticipated price $f_n = \underline{f}$ when shrouding. Since $v - \underline{f} > w - c^w_{\min}$, Step (ii) implies that naive consumers buy the inferior product whenever shrouding occurs. Step (iii) uses Bertrand-competition-type arguments to show that sophisticated consumers always buy a total-surplus-maximizing product at marginal cost. This implies (Step (iv)) that firms can only earn profits if shrouding occurs, and hence all firms shroud with probability 1, a contradiction.

We shall refer to the anticipated price f_n of the inferior product and the total price t^w_n of the superior product as the perceived price (when shrouding) below. Also, if a firm does not offer a contract for the superior product we define $t^w_n = \infty$, and if it does not offer the inferior one we define $f_n = \infty$ and $a_n = 0$.

Step (i): Since shrouding occurs with positive probability, any firm can guarantee itself strictly positive profits through shrouding and offering the contract $(\underline{f}, \overline{a})$ for the inferior product and offering the superior product at or above marginal cost. In this case, the price floor ensures that no rival can offer a contract for the inferior good that naive consumers perceive as being strictly better, and naive consumers buy the superior good only if it is priced below the lowest marginal cost, and such purchase cannot occur with positive probability in equilibrium. Thus, all firms must earn positive expected profits in equilibrium. Let $\pi^*_n > 0$ be firm n's equilibrium expected profits.

Step (ii): Conditional on firm n shrouding, define the perceived deal offered to naive consumers as $\tilde{d}_n = \max\{v - f_n, w - t^w_n\}$. If shrouding occurs, naive consumers choose a best such deal. Let $\underline{\tilde{d}}_n$ be the infimum of $\tilde{d}_n$ conditional on shrouding, and let $\underline{\tilde{d}} = \min_n \{\underline{\tilde{d}}_n\}$. A key step in our proof is the following:

Lemma 2. $\underline{\tilde{d}} = v - \underline{f}$.

B.3.1. Proof of Lemma 2. Suppose, towards a contradiction, that $\underline{\tilde{d}} \neq v - \underline{f}$. Since no firm can set an anticipated price below the floor, $\underline{\tilde{d}} < v - \underline{f}$. We begin by showing that there is a firm n' such that for any $\epsilon > 0$, firm n' has an optimal action that involves shrouding and giving perceived deal $\tilde{d}_{n'}(\epsilon)$ such that $0 \leq \tilde{d}_{n'}(\epsilon) - \underline{\tilde{d}} < \epsilon$ and the expected profits of firm n' when choosing this action and shrouding occurring is less than ϵ. Note that if a firm makes strictly greater expected profit on a sophisticated consumer than on a naive consumer conditional on shrouding, then it strictly prefers to unshroud, a contradiction. Hence, it is sufficient to prove our statement for naive consumers. This is trivial from the definition of $\underline{\tilde{d}}$ if no firm sets $\underline{\tilde{d}}$ with positive probability when shrouding, or exactly one firm sets $\underline{\tilde{d}}$ with positive probability when shrouding. Suppose, therefore, that two or more firms set $\underline{\tilde{d}}$ with positive probability when shrouding. Let n' be one of these firms. We will show that if firm n' earns positive expected profits from this action conditional on shrouding occurring, then it has a profitable deviation. If it earns positive expected profits on the inferior product only, lowering $f_{n'}$ minimally while still shrouding and holding other prices constant is a profitable deviation; if it earns positive expected profits on both the inferior and the superior product, then lowering $f_{n'}$ and $t^w_{n'}$ minimally by the same amount is a profitable deviation; and if it earns positive expected profits on the superior product only, then unshrouding and lowering $t^w_{n'}$ minimally is a profitable deviation. This concludes our argument that a firm n' with the above properties exists. Denote a deal offered from a firm n to sophisticated consumers, or its 'real deal', by $d_n = \max\{w - t^w_n, v - f_n - a_n\}$. For future reference, note two facts: (i) our undercutting argument implies that if firm n' chooses $\underline{\tilde{d}}$ and shrouds, conditional on shrouding occurring it earns zero profits; and (ii) we can choose sequences $\epsilon^l \to 0$ and $\tilde{d}_{n'}(\epsilon^l)$ such that the corresponding real deals converge; denote the limit by $d(\underline{\tilde{d}})$.

For a sufficiently small $\epsilon > 0$, the optimal action in which firm n' sets the perceived deal $\tilde{d}_{n'}(\epsilon)$ earns less than ϵ, so firm n' must earn $\pi^*_{n'} - \epsilon$ conditional on unshrouding occurring. Now for it to be the case that firm n' makes $\pi^*_{n'} - \epsilon$ when unshrouding occurs but at most ϵ when shrouding occurs, there must exist a firm $n'' \neq n'$ and an $\epsilon' > 0$ such that firm n'' offers a strictly better real deal than the real deal associated with $\tilde{d}_{n'}(\epsilon)$ with probability of at least $1 - \epsilon'$ when it shrouds but offers a worse real deal than that associated with $\tilde{d}_{n'}(\epsilon)$ with positive probability when it unshrouds. Furthermore, as $\epsilon \to 0$, we must have $\epsilon' \to 0$, and the probability with which firm n'' offers a worse real deal than that associated with $\tilde{d}_{n'}(\epsilon)$ when it unshrouds remains bounded away from zero. This implies that there is a firm n'' such that firm n'' offers a weakly better real deal than $d(\underline{\tilde{d}})$ with probability 1 when it shrouds but offers a weakly worse real deal than $d(\underline{\tilde{d}})$ with positive probability when it unshrouds.

For each firm n that unshrouds with positive probability, denote the infimum of real deals offered by firm n to consumers conditional on unshrouding by $\underline{d}_n$. Let $\underline{d}$ be the minimum of these real deals. Note that for a sufficiently small $\epsilon > 0$, the real deal associated with $\tilde{d}_{n'}(\epsilon)$ is strictly better than $\underline{d}$ as conditional on unshrouding occurring firm n' must make positive profits when offering the perceived deal $\tilde{d}_{n'}(\epsilon)$. This implies that $d(\underline{\tilde{d}})$ is a weakly better deal than $\underline{d}$.

First, suppose that some firm offers $\underline{d}$ with positive probability when unshrouding. Observe that two firms cannot offer this deal with positive probability, as otherwise a firm would benefit from offering a minimally better deal since it must make positive expected profits when offering $\underline{d}$. We argue that if there is a single firm that offers $\underline{d}$ with positive probability, it must be firm n''. Suppose firm $n''' \neq n''$ offers $\underline{d}$ with positive probability when unshrouding. Then, conditional on shrouding firm n'' cannot offer a real deal equal to $\underline{d}$ with positive probability; otherwise, firm n''' would prefer to undercut. Hence, using that firm n'' offers a weakly better real deal than $d(\underline{\tilde{d}}) \geq \underline{d}$ with probability 1 when it shrouds, we conclude that firm n'' offers a strictly better real deal than $\underline{d}$ with probability 1 when it shrouds. This implies that firm n''' makes zero profits conditional on firm n'' shrouding. In addition, firm n''' makes zero profits conditional on firm n'' unshrouding as $\underline{d} \leq \underline{d}_{n''}$, and both cannot be set with positive probability—a contradiction. This completes the argument that firm n'' must offer $\underline{d}$ with positive probability.

Again using that it cannot be the case that two firms set $\underline{d}$ with positive probability, all rivals of firm n'', including firm n', must with positive probability offer a strictly worse real deal than $\underline{d}$ conditional on shrouding. Choose an optimal such offer, and denote the real deal by $d' < \underline{d}$ and the perceived deal by $\tilde{d}'$. Because $d' < \underline{d}$, when offering this deal firm n' earns zero expected profits conditional on unshrouding occurring, and hence it must earn more than $\pi^*_{n'}$ conditional on shrouding occurring. Using the fact from above that firm n' earns zero profits if it sets $\underline{\tilde{d}}$ and shrouding occurs, this implies that the perceived deal $\tilde{d}'$ must be strictly better than $\underline{\tilde{d}}$. Also when shrouding occurs and firm n' makes this offer, firm n' cannot attract sophisticated consumers because firm n'' offers a better deal than $d(\underline{\tilde{d}}) \geq \underline{d} > d'$ with probability 1 when shrouding. Thus, firm n' attracts only naive consumers when offering $\tilde{d}'$. Notice that firm n' cannot attract only naive consumers to product w as naive consumers evaluate product w in the same way as sophisticated consumers do and evaluate the inferior one as a (weakly) better deal than sophisticated ones do. Hence, when firm n' makes the above offer, naive consumers buy product v from firm n' with positive probability. Now choose another optimal offer by firm n' in which it shrouds, offers a perceived deal $\hat{\tilde{d}} < \tilde{d}'$, makes lower profits conditional on shrouding occurring, and earns positive profits conditional on unshrouding occurring (this is possible because $\tilde{d}' > \underline{\tilde{d}}$). Denote this offer by $\hat{t}^w, (\hat{f}, \hat{a})$. The fact that firm n' earns positive expected profits conditional on unshrouding occurring means that the corresponding real deal $\hat{d}$ must be weakly better than $\underline{d}$, so that $\hat{d} > d'$. Recall that firm n' must with positive probability attract naive consumers to the inferior product when offering $\tilde{d}'$. For the deal $\tilde{d}'$ to be worse but at the same time to be perceived as better than $\hat{\tilde{d}}$, it must be that

$$v - f' > \max\{w - \hat{t}^w, v - \hat{f}\},$$

$$v - f' - a' < \max\{w - \hat{t}^w, v - \hat{f} - \hat{a}\},$$

where (f', a') is a contract offer that corresponds to the perceived deal $\tilde{d}'$. Since firm n' must attract naive consumers with positive probability when offering $\tilde{d}'$, it is strictly better off changing the perceived deal $\hat{\hat{d}}$ to $\tilde{d}'$ while holding the real deal constant. To complete the argument, we show that this deviation is feasible. In case $f'+a' \geq \hat{f}+\hat{a}$, firm n' can do so by lowering $\hat{f}$ to f' and raising the additional price so as to keep the total price of the inferior product fixed. If $f'+a' < \hat{f}+\hat{a}$, then the second inequality above implies that $\max\{w-\hat{t}^w, v-\hat{f}-\hat{a}\} = w-\hat{t}^w$ and hence firm n' can do so by changing $(\hat{f}, \hat{a})$ to (f', a'), thereby improving the perceived deal without affecting the real deal $w-\hat{t}^w$ it offers to sophisticated consumers as well as naive consumers when unshrouding occurs. We thus have a contradiction.

Secondly, suppose that no firm offers $\underline{d}$ with positive probability. Note that in this case $d(\tilde{\underline{d}}) > \underline{d}$; otherwise, for a sufficiently small $\epsilon > 0$ firm n' could not be making $\pi^*_{n'}$ when offering the perceived deal $\tilde{d}_{n'}(\epsilon)$, as it would make lower profits both when shrouding occurs and when unshrouding occurs. By the following argument, which mimics the one above, n'' must achieve the infimum $\underline{d}$. To see it, note that if a firm $n''' \neq n''$ achieves this infimum, then since n'' offers a weakly better deal than $d(\tilde{\underline{d}}) > \underline{d}$ when shrouding, firm n''' earns zero profits conditional on n'' shrouding; and as the real deals of n''' approach $\underline{d}$ the probability of n'' offering a better real deal when unshrouding goes to 1, so n''' expected profits conditional on n'' unshrouding go to zero. This contradicts that firm n''' must earn $\pi^*_{n'''} > 0$ for almost all offers. Hence, n'' achieves the infimum $\underline{d}$.

Take a sequence of real deals $d^l_{n''} \to \underline{d}$ that are optimal for firm n'' when unshrouding. Then, the expected profits firm n'' earns from unshrouding and choosing $d^l_{n''}$ when firm n' unshrouds approach zero as $l \to \infty$. Hence, it must be the case that conditional on shrouding, firm n' charges a weakly worse deal than $\underline{d}$ with positive probability. Choose an optimal such offer, and denote the real deal by $d' \leq \underline{d}$ and the perceived deal by $\tilde{d}'$. From here, the logic is essentially the same as when a single firm charges $\underline{d}$ with positive probability, but we repeat a large part of it because there are minor changes. Because $d' \leq \underline{d}$ and no firm charges $\underline{d}$ with positive probability when unshrouding, when offering this deal firm n' earns zero expected profits conditional on unshrouding occurring, and hence it must earn at least $\pi^*_{n'} > 0$ conditional on shrouding occurring. Using the fact from above that firm n' earns zero profits if it sets $\tilde{\underline{d}}$ and shrouding occurs, this implies that the perceived deal $\tilde{d}'$ must be strictly better than $\tilde{\underline{d}}$. Also when shrouding occurs and firm n' makes this offer, firm n' cannot attract sophisticated consumers because firm n'' offers a weakly better deal than $d(\tilde{\underline{d}})$ with probability 1 when shrouding, and $d(\tilde{\underline{d}}) > \underline{d} \geq d'$. Thus, firm n' attracts only naive consumers when offering $\tilde{d}'$. Notice that firm n' cannot attract only naive consumers to product w as naive consumers evaluate product w in the same way as sophisticated consumers do and evaluate the inferior one as a (weakly) better deal than sophisticated ones do. Hence, when firm n' makes the above offer, naive consumers buy product v from firm n' with positive probability. Now choose another optimal offer by firm n' in which it shrouds, offers a perceived deal $\hat{\hat{d}} < \tilde{d}'$, makes lower profits conditional on shrouding occurring, and earns positive profits conditional on unshrouding occurring (this is possible because $\tilde{d}' > \tilde{\underline{d}}$). Denote the real deal of this offer by $\hat{d}$. The fact that firm n' earns positive expected profits conditional on unshrouding occurring means that $\hat{d}$ must be strictly better than $\underline{d}$, so that $\hat{d} > d'$. From here, the argument is exactly the same as when a single firm charges $\underline{d}$ with positive probability. This completes the proof of the lemma. ‖

Hence, we conclude that $\tilde{\underline{d}} = v - \underline{f}$, and, therefore, all firms set an anticipated price $\underline{f}$ for the inferior product with probability 1 conditional on shrouding and all naive consumers buy the inferior product when shrouding occurs (since $v - \underline{f} > w - c^w_{\min}$ and no firm in equilibrium sells the superior product at a price $t^w_n < c^w_{\min}$ with positive probability).

Step (iii): We next establish that there exists a firm n that offers a deal in which $d_n = w - c^w_{\min}$ with probability 1. Suppose otherwise. Then there exists a real deal d^0 at which a most efficient firm, *i.e.*, a firm with cost $c^w_{\min}$ for the superior product, earns positive expected profits π^0 from selling to sophisticated consumers. Let $\underline{\underline{d}}_n$ be the infimum of the real deals set by firm n when either shrouding or unshrouding. Let $\underline{\underline{d}} = \min_n\{\underline{\underline{d}}_n | c^w_n = c^w_{\min}\}$. First, if multiple firms offer $\underline{\underline{d}}$ with positive probability, a most efficient firm that does so with positive probability when unshrouding must sell at $\underline{\underline{d}}$ with positive probability, and hence prefers to minimally raise $\underline{\underline{d}}$. If a most efficient firm sets $\underline{\underline{d}}$ with positive probability when shrouding, it can minimally raise $\underline{\underline{d}}$ keeping the price for the inferior product $(\underline{f}, a_n)$ as well as its shrouding decision fixed; this increases its profits from sophisticated consumers without affecting the behaviour of naive consumers when shrouding occurs, and at most minimally decreases its profits from selling to naive consumers when unshrouding occurs. Secondly, if a single most efficient firm charges $\underline{\underline{d}}$, it must be shrouding to earn positive profits. But then keeping the price for the inferior product $(\underline{f}, a_n)$ as well as its shrouding decision fixed and at the same time offering the deal d^0 is a profitable deviation, because it does not affect the profits from selling to naive consumers when shrouding occurs, strictly increases the profits from selling to sophisticated consumers, and weakly increases the profits from selling to naive consumers if unshrouding occurs. If no most efficient firm offers the deal $\underline{\underline{d}}$ with positive probability, choose a most efficient firm that achieves this infimum. Consider a sequence d^l of optimal real deals offered by this firm that converges to $\underline{\underline{d}}$. For sufficiently large l, the firm earns less than $\pi^0 > 0$ from sophisticated consumers. Hence, the firm must shroud when making these offers. But then again keeping the price for the inferior product $(\underline{f}, a_n)$ as well as its shrouding decision fixed, and at the same time offering the deal d^0 is a profitable deviation. Hence, we conclude that there exists a firm n that offers a deal in which $d_n = w - c^w_{min}$ with probability 1.

Step (iv): Since no firm can offer more than the maximal total surplus $w-c^w_{\min}$ without making a loss, and only a most efficient firm can make an offer in which $d_n = w-c^w_{\min}$ without making a loss, with probability 1 all sophisticated consumers buy from a most efficient firm. Furthermore, because no firm can earn positive profits if unshrouding occurs, every firm shrouds with probability 1, a contradiction.

We have, therefore, established that shrouding cannot occur with an interior probability. To complete the proof of the proposition, we establish properties of a deceptive equilibrium. Step (i) still applies, so all firms earn positive expected profits. If shrouding occurs with probability 1, then the same steps as in the first paragraph of the proof of Lemma 2 imply that if $\underset{\sim}{d} < v-\underline{f}$, there is a firm and an optimal action that does not earn the firm's equilibrium profits, a contradiction. Hence, again all firms set an anticipated price $\underline{f}$ for the inferior product with probability 1 and all naive consumers buy the inferior product. Now the exact same proof as in Step (iii) above implies that there is a firm that offers a real deal $w-c^w_{\min}$ with probability 1. Furthermore, to earn positive profits, every firm must set a real deal for the inferior product that is strictly worse than $w-c^w_{\min}$. Hence, only naive consumers buy the inferior product. Because every firm must attract some naive consumers to earn positive profits, every firm sets $f_n = \underline{f}$, and, since firm n's market share is independent of a_n, every firm sets $a_n = \overline{a}$. Hence, any deceptive equilibrium must have the properties in the proposition. Finally, if at least two most efficient firms charge $c^w_{\min}$ for product w, then obviously no firm wants to deviate, so that a deceptive equilibrium indeed exists. This completes the proof of the proposition. ‖

B.4. *Proof of Proposition 4*

Suppose, towards a contradiction, that such a deceptive equilibrium exists. Since naive consumers ignore the additional price in their purchase decision and firms earn zero profits from selling to sophisticated consumers, every firm with a positive market share for product v must set $a_n = \overline{a}$ to maximize profits. Furthermore, every firm can guarantee positive profits from selling to naive consumers by offering product v at $(\underline{f}, \overline{a})$. Let $\overline{f}_n$ be the supremum of firm n's anticipated price distribution for v, and let $\overline{f} = \max_n\{\overline{f}_n\}$. Because firms earn zero profits from selling to sophisticated consumers who buy w, the same arguments as in the proof of Proposition 2 imply that $\overline{f} = \underline{f}$. Hence, in any deceptive equilibrium all firms set the $f_n = \underline{f}, a_n = \overline{a}$. Since sophisticated consumers buy product w, it must be the case that $v-(\underline{f}+\overline{a}) \le w-c^w_{\min}$. Hence, firm n can attract all consumers by unshrouding and setting its total price t_n such that $v-t_n = w-c^w_{\min}+\epsilon'$ for any $\epsilon' > 0$. This means that firm n earns a profit of $t_n - c_n = (v-c_n)-(w-c^w_{\min})-\epsilon' > \epsilon-\epsilon'$, so that firm n can achieve a profit arbitrarily close to ϵ. Since industry profits are bounded, for a sufficiently large N there must be a firm for which this is a profitable deviation. ‖

B.5. *Proof of Proposition 5*

We construct such a deceptive equilibrium for the case in which $v-c_{\min} \le w-c^w_{\min}$. Note that the assumptions $v > w$ and $\underline{f} \le c^w_{\min}$ imply $v-\underline{f} > w-c^w_{\min}$. Consider a candidate equilibrium in which all firms shroud, sell v at $(\underline{f}, \overline{a})$, and w at $(c^w_n, 0)$ with probability 1. In the candidate equilibrium, sophisticated consumers buy w from a firm selling it at the lowest possible cost $c^w_{\min}$ because $v-\underline{f}-\overline{a} < v-c_{\min} \le w-c^w_{\min}$. Naive consumers buy v because $v-\underline{f} > w-c^w_{\min}$, and all firms earn positive profits from selling to naive consumers since $\underline{f}+\overline{a} > c_n$ for all n. Any deviation involving unshrouding is unprofitable because a competitor offers the superior product at the lowest possible costs, and thus consumers understanding the total price cannot be profitably attracted. Similarly, conditional on shrouding it is impossible to profitably attract sophisticated consumers, or to attract more naive consumers since all competitors offer the highest possible perceived deal $v-\underline{f}$. Raising the anticipated price for product v reduces the demand from naive consumers and thus profits to zero; lowering the additional price for product v leaves demand unaffected, and hence lowers profits. Repricing w cannot attract naive consumers since $v-\underline{f} > w-\underline{f}$, and hence we have constructed a deceptive equilibrium with the desired features.

Define $t^w_n = f^w_n + a^w_n$. Because firms earn zero profits from selling to sophisticated consumers in any such equilibrium, sophisticated consumers buy w at $t^w_n = c^w_{\min}$. Given this, the proof of Proposition 4 applies unaltered and establishes that such a deceptive equilibrium does not exist if v is non-vanishingly superior to product w. ‖

B.6. *Proof of Proposition 6*

We first state a deceptive equilibrium satisfying the properties listed in the proposition for the case in which $M \le s_n(\underline{f}+\overline{a}-c_n)$, and verify that it is indeed an equilibrium. Let all firms shroud and charge $(\underline{f}, \overline{a})$ for product v. Let all firms $n' \neq n$ charge $c^w_{n'}$ for the transparent product w, and let firm n charge $c^w_{\min}$ for product w. Finally, naive consumers buy product v and sophisticated consumers buy product w from firm n when facing the equilibrium prices.

Since $v-\underline{f} > w-c_n^w > w-c_{\min}^w$, naive consumers strictly prefer buying v. Furthermore, $w-c_{\min}^w > v-c_{\min}$ and $\underline{f}+\overline{a} > c_{\min}$ implies that $w-c_{\min}^w > v-(\underline{f}+\overline{a})$, so sophisticated consumers prefer buying w, and hence it is optimal for them to buy w from firm n.

Conditional on shrouding and at most one firm deviating, naive consumers buy the inferior product as long as no firm prices w below c_n^w since $v-\underline{f} > w-c_n^w$, and sophisticated consumers buy product w as long as no firm charges v below $c_{\min}$ since $w-c_{\min}^w > v-c_{\min}$. Using these facts, it is straightforward to verify that conditional on shrouding charging $(\underline{f},\overline{a})$ for v is optimal, and that the pricing for w is also optimal. Similarly, conditional on unshrouding it is impossible to attract consumers to product v while pricing weakly above costs, and using this fact it is also straightforward to verify that the specified pricing remains optimal conditional on unshrouding. If any firm other than n sets the same prices and unshrouds, it makes zero profits, so it does not unshroud. Finally, firm n makes a profit of M on every sophisticated consumer whether or not it unshrouds, so its unshrouding decision hinges on profits from naive consumers. If firm n shrouds, it makes $s_n(\underline{f}+\overline{a}-c_n)$ on the pool of naive consumers, and if it unshrouds, it can earn at most M per consumer. Hence, shrouding by firm n is optimal.

We next prove the converse: that if $M > s_n(\underline{f}+\overline{a}-c_n)$, a deceptive equilibrium with the properties stated in the proposition does not exist. Because firm n sells product w at price $c_{\min}^w$, any firm $n' \neq n$ sets the price of product w at or above $c_{\min}^w$; otherwise, sophisticated consumers do not buy from firm n. If $M > s_n(\underline{f}+\overline{a}-c_n)$, then firm n strictly prefers to deviate from the candidate deceptive equilibrium in the proposition by unshrouding and selling product w to all consumers at price $c_{\min}^w-\epsilon$ for sufficiently small $\epsilon > 0$, so such an equilibrium does not exist. ‖

B.6.1. Proof of Corollary 2. Since $\underline{f} \leq \tilde{v}-\underline{u}$ and $\overline{a} \leq v-\tilde{v}$, we have $\underline{f}+\overline{a} \leq v-\underline{u}$, which implies that condition (SC) is violated. ‖

B.7. *Proof of Lemma 1*

We first show that the sets of Nash-equilibrium outcomes when f_n is unrestricted and when f_n is restricted to be nonnegative are the same. Suppose, first, that unshrouding occurs with probability 1 and $v-\underline{u} < c_{\min}$. Then, in both the restricted and unrestricted games consumers buy with probability 0, so the outcomes are the same.

To prove the first statement in the lemma otherwise, we begin by arguing that any equilibrium in the restricted game is an equilibrium in the unrestricted game by establishing that firm n does not strictly prefer to charge $f_n < 0$ when all other firms set nonnegative anticipated prices. We start with some preliminary observations regarding the consumers' behaviour. Note that a consumer's charging decision depends only on the cards she accepts, the additional prices of these cards, and whether unshrouding occurs (and, conditional on the set of cards she accepts, not on anticipated prices). Note also that if all firms shroud, consumers accept any card for which $f_n < 0$ and charge their favourite card n' among those they accept. When it comes to period 2, a consumer who charged her expenses to card n' delays repayment to period 3 if $\beta R_{n'} < 1$, or $\beta < 1/R_{n'}$. As a result, firm n' earns an expected *ex post* profit of $Q(1/R_{n'})(R_{n'}-1)$ from the consumers who charged their expenses to card n'. If unshrouding occurs, consumers accept all cards for which $f_n < 0$ if they prefer to accept some card, *i.e.*, if $\sum_n I_{\{f_n<0\}} f_n + \min_{n'}\{I_{\{f_{n'}>0\}} f_{n'} + Q(1/R_{n'})(R_{n'}-1)\} \leq v-\underline{u}$, where $I_{\{\}}$ denotes the indicator function. And if unshrouding occurs and consumers get a card, they charge their expenses to their preferred card among those that minimize $I_{\{f_{n'}>0\}} f_{n'} + Q(1/R_{n'})(R_{n'}-1)$.

Now to see why firm n does not strictly prefer to set $f_n < 0$, let firm n set $f_n = 0$ instead of $f_n < 0$ while leaving its total price $f_n + Q(1/R_n)(R_n-1)$ and its unshrouding decision the same. Any consumer who gets and charges firm n's card with the former prices also does so with the latter prices, while with the latter prices no other consumer gets firm n's card. Hence, the latter prices cannot be less profitable.

Next, we prove that any equilibrium in the unrestricted game is an equilibrium in the restricted game by proving that no firm charges $f_n < 0$ with positive probability in any Nash equilibrium of the unrestricted game. It follows from our observations about consumer behaviour above that if two or more firms set $f_n < 0$, each consumer either accepts all of these cards or does not accept any card. We consider the following two cases.

First, suppose shrouding occurs with positive probability in equilibrium. Then, every firm can ensure positive profits by shrouding and setting $f_n = 0$ and $R_n = \operatorname{argmax}_R Q(1/R)(R-1)$, which will lead all consumers whose favourite card is card n to get it and charge it whenever shrouding occurs. Hence, all firms earn positive profits. Notice that if there exist some consumers who accept card n' with $f_{n'} < 0$ but do not charge their expenses to card n', then firm n' can profitably deviate by raising $f_{n'}$ by a bit and lowering $R_{n'}$ while keeping $f_{n'} + Q(1/R_{n'})(R_{n'}-1)$ constant—this ensures that both when shrouding occurs and when unshrouding occurs firm n' loses no consumers who accept and charge card n', and firm n' increases profits from consumers who accept but do not charge card n'. Hence, if some consumers accept card n' with $f_{n'} < 0$, all of the consumers charge their expenses to card n'.

We now show that at most one firm n' can set a negative anticipated price. To see why, suppose otherwise: firm n' and n'' set negative anticipated prices with positive probability. Since all consumers who accept card n' with $f_{n'} < 0$ charge their expenses to card n' and since each consumer either accepts all cards with negative anticipated prices or does not accept any card, all consumers must accept *no* card with probability 1 whenever two firms set negative anticipated prices. This requires that unshrouding occurs with probability 1 whenever both firms set negative anticipated prices, and hence one firm, say n', must unshroud with probability 1 whenever it sets a negative anticipated price $f_{n'}<0$. Since conditional on $f_{n'}<0$ consumers do not accept card n' whenever the other firm n'' also sets $f_{n''}<0$, card n' is not worth accepting for any consumer: $f_{n'}+Q(1/R_{n'})(R_{n'}-1)>v-\underline{u}$ whenever $f_{n'}<0$. But then, firm n' earns zero profits when setting $f_{n'}<0$, a contradiction. Hence, at most one firm, say firm n', sets a negative anticipated price when shrouding occurs with positive probability.

Take any firm $n \neq n'$ that does not set a negative anticipated price. Because all consumers accept only card n' whenever $f_{n'}<0$, firm n must earn positive expected profits conditional on all firms setting a nonnegative anticipated price. To do so, firm n cannot with positive probability set $f_n>v-\underline{u}$ when it shrouds because with such an anticipated price it will make zero sales both when shrouding occurs and when unshrouding occurs and all firms set nonnegative anticipated prices. Furthermore, firm n cannot with positive probability set $f_n+Q(1/R_n)(R_n-1)>v-\underline{u}$ when it unshrouds as otherwise firm n makes zero sales when all firms set nonnegative anticipated prices. Similarly, firm n' must with probability 1 set $f_{n'}\leq v-\underline{u}$ conditional on shrouding and $f_{n'}+Q(1/R_{n'})(R_{n'}-1)\leq v-\underline{u}$ conditional on unshrouding. This implies that all consumers with probability 1 accept a card both when shrouding occurs and when unshrouding occurs, and, therefore, consumers always accept any card that charges a negative anticipated price. Thus, firm n' that charges a negative anticipated price $f_{n'}<0$ with positive probability can gain by replacing each negative anticipated price $f_{n'}$ through a higher anticipated price $f_{n'}/2$ and keeping its strategy otherwise unchanged. Hence, in any equilibrium all firms charge nonnegative prices with probability 1 if shrouding occurs with positive probability.

Secondly, suppose unshrouding occurs with probability 1 and $v-\underline{u}>c_{\min}$. It is easy to see that all consumers must accept some card with probability 1 in this case. As above, all consumers who accept card n with $f_n<0$ must charge their expenses to card n, and hence with probability 1 firm n must set the lowest total price whenever setting $f_n<0$. Using this fact, standard Bertrand-type arguments applied to the total price $f_n+Q(1/R_n)(R_n-1)$ establish that in equilibrium at least two most efficient firms must charge a total price $f_n+Q(1/R_n)(R_n-1)=c_{\min}$ with probability 1, no firm charges a total price strictly below $c_{\min}$ with positive probability, and consumers buy with probability 1. But then any firm n that charges $f_n<0$ attracts some consumers that do not charge their expenditures with the firm; firm n would be strictly better off when charging the same total expected payment and setting $f_n\geq 0$. This completes the argument that no firm sets a negative anticipated price in any Nash equilibrium of the unrestricted game, and thereby the proof that the sets of Nash-equilibrium outcomes in the restricted and unrestricted games are the same.

We now move on to the second part of the lemma. Given that $f_n\geq 0$ for all n, each consumer gets at most one card, and if she does charges it. As a result, firm n earns an expected *ex post* profit of $a_n\equiv Q(1/R_n)(R_n-1)$ from all its consumers. The maximum *ex post* profit per such consumer the firm can earn is $\overline{a}=\max_R Q(1/R)(R-1)$, which exists by assumption. Notice that (i) firm n's profit from a consumer is f_n+a_n; (ii) a consumer obtains utility $v-f_n-a_n$ from card n; (iii) in period 0, consumers expect to receive utility $v-f_n$ from card n if shrouding occurs and $v-f_n-a_n$ if unshrouding occurs. Hence, thinking of the firms as choosing $a_n\in[0,\overline{a}]$ rather than R_n generates the same payoffs as in our reduced-form model, making the games strategically equivalent. Finally, note that consumers' lexicographic preferences over cards induce exogenously given market shares when consumers are indifferent between cards. ‖

B.8. *Proof of Proposition 7*

Recall that a deceptive equilibrium exists for $\eta=0$ if and only if for all firms n,

$$s_n(\underline{f}+\overline{a}-c_n)\geq v-\underline{u}-c_n.$$

This proof has five steps.

Step (i): No firm unshrouds the additional price with probability 1. If a firm unshrouds with probability 1, all consumers become sophisticated and hence buy from the firm with the lowest total price $f+a$. Hence by standard Bertrand arguments, all consumers buy at a total price $f+a=c_{\min}$ and no firm makes positive profits from selling to the consumers excluding the unshrouded cost. Then, the firm that chooses to unshroud makes negative profits—a contradiction.

Step (ii): All firms earn positive profits. According to (i), in any equilibrium there is positive probability that no firm unshrouds. Then, each firm n can earn positive profits by shrouding the additional prices and offering $(\underline{f},\overline{a})$.

Step (iii): The distributions of total prices are bounded from above. Suppose firm n sets the total price $f_n+a_n>v-\underline{u}+\overline{a}$ with positive probability in equilibrium. When the additional prices are shrouded, consumers never buy the product from firm n because this inequality implies $f_n>v-\underline{u}$. When the additional prices are unshrouded, consumers never buy from firm n because $f_n+a_n>v-\underline{u}$. Firm n's profits in this case are at most zero, a contradiction with (ii).

Step (iv): No firm unshrouds the additional price with positive probability. Let $\hat{t}_n$ be the supremum of the equilibrium total-price distribution of firm n when unshrouding; set $\hat{t}_n=0$ in case firm n does not unshroud. Let $\hat{t}=\max_n\{\hat{t}_n\}$; by (iii),

$\hat{t}$ is bounded from above and hence well defined. Consider firm n that unshrouds and for whom $\hat{t}_n = \hat{t}$. Note that in any equilibrium in which some firm unshrouds with positive probability, $\hat{t} > c_n$ by (ii).

First, suppose that firm n charges the total price $\hat{t}$ with positive probability. If some other firm $n' \neq n$ also sets the total price $\hat{t}$ with positive probability, then firm n has an incentive to slightly decrease its total price—a contradiction. Thus, only firm n charges the total price $\hat{t}$ with positive probability. Because $\hat{t}$ is the supremum of the total-price distribution conditional on unshrouding, firm n can earn positive profits only if all firms other than n choose to shroud. Conditional on all other firms shrouding, n's expected profits are no larger than $v - \underline{u} - c_n - \eta$, because the additional price is unshrouded by firm n and hence consumers never buy the product from firm n if $\hat{t} > v - \underline{u}$. When firm n shrouds and offers $(\underline{f}, \overline{a})$, however, its profits conditional on all other firms shrouding are at least $s_n(\underline{f} + \overline{a} - c_n)$. Thus, the equilibrium condition $s_n(\underline{f} + \overline{a} - c_n) \geq v - \underline{u} - c_n$ implies that deviating by shrouding and offering $(\underline{f}, \overline{a})$ is profitable—a contradiction.

Secondly, suppose that firm n does not charge the total price $\hat{t}$ with positive probability. Then, for any $\epsilon > 0$, firm n charges a total price in the interval $(\hat{t} - \epsilon, \hat{t})$ with positive probability. As $\epsilon \to 0$, the probability that firm n conditional on some other firm unshrouding can attract consumers goes to zero, because $\hat{t}$ is the supremum of the total-price distribution conditional on unshrouding. Hence, firm n cannot earn the unshrouding cost $\eta > 0$ conditional on some other firm unshrouding—*i.e.*, it loses money in expectation relative to shrouding and offering $(\underline{f}, \overline{a})$. In addition, conditional on all other firms shrouding firm n earns less than $v - \underline{u} - c_n$: the deviation profits in the no-unshrouding-cost case. Because shrouding is an equilibrium in the no-unshrouding-cost case, there is a profitable deviation for firm n—a contradiction.

Step (v): All firms offer the contract $(\underline{f}, \overline{a})$ with probability 1. By (iv), all firms choose to shroud with probability 1. By the exact same argument as in Proposition 2 for the case in which shrouding occurs with probability 1, all firms offer the contract $(\underline{f}, \overline{a})$ with probability 1. ‖

Acknowledgments. Kőszegi thanks the Berkeley Center for Equitable Growth, and Heidhues and Kőszegi thank the European Research Council (Starting Grant #313341) for financial support. We are grateful to Özlem Bedre-Defolie, Stefano DellaVigna, Drew Fudenberg, Michael Grajek, Michael Grubb, Péter Kondor, Nicolas Melissas, Klaus Schmidt, Antoinette Schoar, Cédric Wasser, and seminar and conference audiences for useful comments and feedback.

REFERENCES

AGARWAL, S. CHOMSISENGPHET, S. MAHONEY, N. *et al*. (2015), "Regulating Consumer Financial Products: Evidence from Credit Cards", *Quarterly Journal of Economics*, **130**, 111–164. Previously NBER Working Paper No. 19484.

AGARWAL, S. DRISCOLL, J. C. GABAIX, X. *et al*. (2009), "Financial Decisions over the Life-Cycle and Implications for Regulation", *Brookings Papers on Economic Activity*, **2009**, 51–117.

AGARWAL, S. and EVANOFF, D. D. (2013), "Loan Product Steering in Mortgage Market" (Working Paper No. 2204400, SSRN).

ALEXANDER, G. J., JONES, J. D. and NIGRO, P. J. (1998), "Mutual Fund Shareholders: Characteristics, Investor Knowledge, and Sources of Information", *Financial Services Review*, **7**, 301–316.

ANAGOL, S. and KIM, H. (2012), "The Impact of Shrouded Fees: Evidence from a Natural Experiment in the Indian Mutual Funds Market", *American Economic Review*, **102**, 576–593.

ARMSTRONG, M. and VICKERS, J. (2012), "Consumer Protection and Contingent Charges", *Journal of Economic Literature*, **50**, 477–493.

AUGENBLICK, N., NIEDERLE, M. and SPRENGER, C. (2015), "Working over Time: Dynamic Inconsistency in Real Effort Tasks", *The Quarterly Journal of Economics*, **130**, 1067–1115.

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

BARBER, B. M., ODEAN, T. and ZHENG, L. (2005), "Out of Sight. Out of Mind: The Effects of Expenses on Mutual Fund Flows", *Journal of Business*, **78**, 2095–2119.

BAR-GILL, O. (2004), "Seduction by Plastic", *Northwestern University Law Review*, **98**, 1373–1434.

BAR-GILL, O. (2008), "The Law, Economics and Psychology of Subprime Mortgage Contracts", (Law and Economics Research Paper No. 08-59, NYU).

BAR-GILL, O. and BUBB, R. (2012), "Credit Card Pricing: The CARD Act and Beyond", *Cornell Law Review*, **97.**

BERGSTRESSER, D., CHALMERS, J. M. R. and Tufano, P. (2009), "Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry", *Review of Financial Studies*, **22**, 4129–4156.

BERK, J. B. and GREEN, R. C. (2004), "Mutual Fund Flows and Performance in Rational Markets", *Journal of Political Economy*, **112**, 1269–1295.

BERK, J. B. and van BINSBERGEN, J. H. (2012), "Measuring Managerial Skill in the Mutual Fund Industry", (Working Paper No. 18184, NBER).

BORENSTEIN, S., MACKIE-MASON, J. K. and NETZ, J. S. (2000), "Exercising Market Power in Proprietary Aftermarkets", *Journal of Economics & Management Strategy*, **9**, 157–188.

BUCKS, B. and PENCE, K. (2008), "Do Borrowers Know Their Mortgage Terms?", *Journal of Urban Economics*, **64**, 218–233.

CARHART, M. M. (1997), "On Persistence in Mutual Fund Performance", *Journal of Finance*, **52**, 57–82.
CHETTY, R., LOONEY, A. and KROFT, K. (2009), "Salience and Taxation: Theory and Evidence", *American Economic Review*, **99**, 1145–1177.
CHOI, J. J., LAIBSON, D. and MADRIAN, B. C. (2010), "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds", *Review of Financial Studies*, **23**, 1405–1432.
COCCO, J. F. (2013), "Evidence on the Benefits of Alternative Mortgage Products", *Journal of Finance*, **68**, 1663–1690.
CONGDON, W. J., KLING, J. R. and MULLAINATHAN, S. (2011), *Policy and Choice: Public Finance Through the Lens of Behavioral Economics*, (Washington, DC: Brookings Institution Press).
CRUICKSHANK, D. (2000), "Competition in U.K. Banking: A Report to the Chancellor of the Exchequer", Technical Report, HM Treasury 2000.
DELLAVIGNA, S. and MALMENDIER, U. (2004), "Contract Design and Self-Control: Theory and Evidence", *Quarterly Journal of Economics*, **119**, 353–402.
DUARTE, F. and HASTINGS, J. S. (2012), "Fettered Consumers and Sophisticated Firms: Evidence from Mexico's Privatized Social Security Market", (Working Paper No. 18582, NBER).
ELIAZ, K. and SPIEGLER, R. (2006), "Contracting with Diversely Naive Agents", *Review of Economic Studies*, **73**, 689–714.
ELLISON, G., (2005), "A Model of Add-On Pricing", *Quarterly Journal of Economics*, **120**, 585–637.
ELLISON, G. and ELLISON, S. F. (2009), "Search, Obfuscation, and Price Elasticities on the Internet", *Econometrica*, **77**, 427–452.
EVANS, D. S. and SCHMALENSEE, R. (2005), *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, 2nd edn. (Cambridge and London: MIT Press).
FARRELL, J. and KLEMPERER, P. (2007), "Coordination and Lock-In: Competition with Switching Costs and Network Effects", in Armstrong, M. and Porter, R. (eds) *Handbook of Industrial Organization*, Vol. 3 (Amsterdam North-Holland) 1967–2072.
FRENCH, K. R. (2008), "Presidential Address: The Cost of Active Investing", *Journal of Finance*, **63**, 1537–1573.
GABAIX, X. and LAIBSON, D. (2006), "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets", *Quarterly Journal of Economics*, **121**, 505–540.
GENNAIOLI, N., SHLEIFER, A. and VISHNY, R. W. (2015), "Money Doctors", *Journal of Finance*, **70**, 91–114.
GERARDI, K., GOETTE, L. and MEIER, S. (2009), "Financial Literacy and Mortgage Outcomes". Paper presented at the ASSA meetings, San Francisco, CA, USA.
GRUBB, M. D. (2009), "Selling to Overconfident Consumers", *American Economic Review*, **99**, 1770–1805.
GRUBB, M. D. (2015), "Consumer Inattention and Bill-Shock Regulation", *Review of Economic Studies*, **82**, 219–257.
GRUBER, M. J. (1996), "Another Puzzle: The Growth in Actively Managed Mutual Funds", *Journal of Finance*, **51**, 783–810.
GURUN, U. G., MATVOS, G. and SERU, A. "Advertising Expensive Mortgages", *Journal of Finance*, forthcoming.
HACKETHAL, A., INDERST, R. and MEYER, S. (2010), "Trading on Advice", (CEPR Working Paper No. 8091, London).
HALL, R. E. (1997), "The Inkjet Aftermarket: An Economic Analysis", Unpublished Manuscript, Stanford University.
HEIDHUES, P. and KŐSZEGI, B. (2012), "Exploiting Naivete about Self-Control in the Credit Market", *American Economic Review*, **100**, 2279–2303.
HEIDHUES, P., KŐSZEGI, B. and MUROOKA, T. (2012) "Deception and Consumer Protection in Competitive Markets", in *The Pros and Cons of Consumer Protection* (Stockholm: Swedish Competition Authority) 44–76.
HEIDHUES, P., KŐSZEGI, B. and MUROOKA, T. (2012), "Inferior Products and Profitable Deception" (Working Paper, European School of Management and Technology) https://www.esmt.org/sites/default/files/digital-measures/Inferior%20products-1.pdf.
HEIDHUES, P., KŐSZEGI, B. and MUROOKA, T. (2016), "Exploitative Innovation", *American Economic Journal: Microconomics*, **8**, 1–23.
KO, K. J. (2012), "Disclosure and Price Regulation in a Market with Potentially Shrouded Costs", (Mimeo).
KOSFELD, M. and SCHÜWER, U. (2011), "Add-on Pricing, Naive Consumers, and the Hidden Welfare Costs of Education", (Working Paper No. 8636, CEPR).
KOSOWSKI, R., TIMMERMANN, A. WERMERS, R. and WHITE, H. (2006), "Can Mutual Fund "Stars" Really Pick Stocks? New Evidence from a Bootstrap Analysis", *Journal of Finance*, **61**, 2551–2595.
LAIBSON, D. I. and YARIV, L. (2007), "Safety in Markets: An Impossibility Theorem for Dutch Books" (Mimeo, Caltech).
LAIBSON, D. I. and YARIV, L. REPETTO, A. and TOBACMAN, J. (2007), "Estimating Discount Functions with Consumption Choices over the Lifecycle" (Working Paper No. 13314, NBER).
LAL, R. and MATUTES, C. (1994), "Retail Pricing and Advertising Strategies", *The Journal of Business*, **67**, 345–70.
MEIER, S. and SPRENGER, C. (2010), "Present-Biased Preferences and Credit Card Borrowing", *American Economic Journal: Applied Economics*, **2**, 193–210.
MULLAINATHAN, S., NÖTH, M. and SCHOAR, A. (2011), "The Market for Financial Advice: An Audit Study." (Mimeo, Massachusetts Institute of Technology).
MUROOKA, T. (2013), "Deception under Competitive Intermediation" (Mimeo, UC Berkeley).
PICCIONE, M. and SPIEGLER, R. (2012), "Price Competition under Limited Comparability", *Quarterly Journal of Economics*, **127**), 197–135.

SHAPIRO, C. (1995), "Aftermarkets and Consumer Welfare: Making Sense of Kodak", *Antitrust Law Journal*, **63**, 483–511.

SHILONY, Y. (1977), "Mixed Pricing in Oligopoly", *Journal of Economic Theory*, **14**, 373–388.

SKIBA, P. M. and TOBACMAN, J. (2008), "Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing, Repayment, and Default" (Vanderbilt Law and Economics Research Paper No. 08-33).

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

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

VARIAN, H. R. (1980), "A Model of Sales", *American Economic Review*, **70**, 651–659.

WILCOX, R. T. (2003), "Bargain Hunting or Star Gazing? Investors' Preferences for Stock Mutual Funds", *The Journal of Business*, **76**, 645–663.

WOODWARD, S. E. and HALL, R. E. (2012), "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence", *American Economic Review*, **100**, 511–515.