Inferior Products and Profitable Deception*

Paul Heidhues
ESMT, CEPR and CESifo

Botond Kőszegi
Central European University, UC Berkeley, and CESifo

Takeshi Murooka
UC Berkeley

November 2, 2012

Abstract

We analyze conditions facilitating profitable deception in a simple model of a competitive retail market. Firms selling homogenous products simultaneously set a transparent up-front price and an additional price, and decide whether to unshroud the additional price to naive consumers. To model especially financial products such as banking services, credit cards, and mutual funds, we assume that there is a binding floor on the product's up-front price. Our main results establish that "bad" products—ones that should not even be produced—tend to be more reliably profitable than "good" products. Specifically, (1) in a market with a single socially valuable product and sufficiently many firms, at least one firm is willing to unshroud, so a deceptive equilibrium does not exist and firms make zero profits. But perversely, (2) if the product is socially wasteful, a firm cannot profitably sell a transparent product, so there is no incentive to unshroud and hence 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. **JEL Codes: D14, D18, D21**

*The setup of this paper was developed in our earlier working paper "The Market for Deceptive Products," which we split into two papers. The companion paper is titled "Exploitative Innovation." Kőszegi thanks the Berkeley Center for Equitable Growth for financial support. We are grateful to Özlem Bedre-Defolie, Stefano DellaVigna, Drew Fudenberg, Michael Grajek, Michael Grubb, Péter Kondor, Nicolas Melissas, Antoinette Schoar, and seminar and conference audiences for useful comments and feedback.

1 Introduction

In this paper, we investigate circumstances under which firms sell products by deceiving some consumers about the products' full cost, 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 surplus than the best alternative facilitate deception precisely because they would not survive in the market if consumers understood hidden fees, and therefore firms often make profits on exactly such bad products but not on good products.

Section 2 introduces our model, in which firms are engaged in simultaneous-move price competition to sell a single homogenous product. Building on the seminal model of Gabaix and Laibson (2006), we assume that each firm charges a transparent up-front price as well as an additional price, and unless at least one firm decides to (costlessly) unshroud the additional prices, naive consumers ignore these prices when making purchase decisions. To capture the notion that in some markets, such as banking services, credit cards, and mutual funds, firms cannot return all profits from later charges by lowering initial charges, we deviate from most existing work and posit that there is a floor on the up-front price.[2] We investigate conditions under which a profitable deceptive equilibrium—wherein all firms shroud additional prices—exists. Whenever such an equilibrium exists, it is the most plausible one: it is then the unique equilibrium in the variant of our model in

[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 identified and reached. Investigating trade and portfolio data from a large German bank, for example, Hackethal, Inderst and Meyer (2010) document that "bank revenues from security transactions amount to €2,560 per customer per year" (2.4 percent 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-percent 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. These observations, however, do not mean that the net economic surplus taking all operating costs into account are large or even positive in these markets: for example, fixed entry costs can dissipate any profits from the later stage of serving consumers.

[2] In Heidhues, Kőszegi and Murooka (2011), we provide a microfoundation for the price floor based on the presence of "arbitrageurs" who would take advantage of overly low prices. This microfoundation is an extreme variant of Ellison's (2005) insight (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. We briefly discuss these and other possible reasons for the price floor in Section 2.2 below. Ko (2011), Grubb (2012), and Armstrong and Vickers (2012) also analyze models with variants of our price-floor assumption.

which unshrouding has any cost (no matter how small the cost is), and all firms prefer it over an unshrouded-prices and hence zero-profit equilibrium.

Section 3 presents our basic results. As a benchmark case, we show that if the price floor is binding, only a *zero-profit* deceptive equilibrium exists, and argue that deception may have economically less important consequences in this situation than in our main cases below. We also note that if consumers are sophisticated in that they observe and take into account additional prices, firms can neither deceive nor earn positive profits from consumers. If the price floor is binding and consumers are naive, however, profitable deception may occur. If other firms shroud and the up-front price is at the floor, a firm cannot compete on the up-front price and can compete on the total 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, an equilibrium with profitable deception exists; and we establish that if there is a firm for which this is not the case, in equilibrium additional prices are unshrouded with probability one, 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 to consumers is lower than its production cost), 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 capturing the entire market, so if there is such a firm only a non-deceptive, zero-profit equilibrium exists. Hence, because in an industry with many firms some firm earns low profits, entry into socially valuable industries makes these industries more transparent; and whenever deceptive practices survive in an industry with many firms, our model says that the industry is socially wasteful. Furthermore, our theory suggests a competition-impairing force in socially valuable industries that is likely to have many implications beyond the current paper: because firms face the threat that a low-profit competitor unshrouds in a valuable but not in a wasteful industry, in the former but not in the latter industry they want to make sure competitors earn sufficient profits to maintain profitable shrouding.

In Section 4, we extend our model by assuming that there are both sophisticated and naive consumers in the market. While we find that in our basic single-product market sophisticated consumers can create an incentive to unshroud, we also show that the situation can be radically different in a multi-product market. In particular, if there is a superior and an inferior product, often sophisticated and naive consumers self-separate into buying the former and the latter product, respectively, and sophisticated consumers exert no pressure to unshroud the inferior product's additional price. 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: all it takes for profitable deception to occur in a competitive industry is the existence of an *inferior* product with a shroudable price component and a binding floor on the up-front price, and firms' profits derive entirely from selling this inferior product.

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, it may have an incentive to unshroud to attract naive consumers to itself, especially if—similarly to for instance Vanguard in the mutual-fund market—it sells mostly the superior product. But if the firm's market power is limited and unshrouding is costly, the extent to which the firm educates consumers is also limited. We also consider a specification of consumer naivete in which consumers know all prices but underestimate their own willingness to pay for an add-on, and show that this alternative generates insights similar to those of our basic model.

In Section 6, we discuss the behavioral-economics and classical literatures most closely related to our paper. Our result that firms sell profitable inferior products to unknowing consumers may look reminiscent of a similar potential implication in classical asymmetric-information models. But while in a rational asymmetric-information setting a lower-quality inferior product may be more profitable to sell than a superior product because it is cheaper to produce and consumers do not know its value, our theory predicts that such a product may be more profitable even if it is more expensive to produce and consumers do know its value. In addition, while a growing theoretical literature investigates how firms exploit naive consumers by charging hidden or unexpected fees, previous work has not identified the central role of wasteful and inferior products in maintaining deception and generating profits. Indeed, in most previous models competition returns all of the profits from hidden fees to consumers, so that these models cannot investigate market conditions that facilitate profitable deception. We conclude in Section 7 with mentioning some policy implications of our findings, and by pointing out important further questions raised by our model.

2 Basic Model

2.1 Setup

In this section, we introduce our basic model of a market for potentially deceptive products. $N \geq 2$ firms compete for naive consumers who value each firm's product at $v > 0$ and are looking to buy at most one item. Firms simultaneously set up-front prices f_n and additional prices a_n, and decide whether to costlessly unshroud the additional prices. For simplicity, we assume that a consumer who buys must pay both prices—she cannot avoid the additional price. If all firms shroud, consumers make purchase decisions believing that the total price of product n is f_n. If at least one firm unshrouds, all firms' additional prices become known to all consumers, and consumers make purchase decisions based on the true total prices $f_n + a_n$. We assume that the highest possible additional price firms can impose is $\overline{a} > 0$. If consumers weakly prefer buying and are indifferent between a subset of firms, these firms split the market in proportion to exogenously given shares $s_n \in [0, 1)$.[3]

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 there are at least two firms whose cost is equal to c_{min}. In addition, we assume that $v + \overline{a} > c_n$ for all firms n; a firm with $v + \overline{a} < c_n$ cannot profitably sell its product, so without loss of generality we can

[3] In later sections, we consider several alternatives to and extensions of the above basic framework. In Section 4, we add sophisticated consumers, and also analyze multiproduct markets. In Section 5, we consider what happens when unshrouding is costly; discuss an alternative formulation of consumer naivete in which there is an add-on consumers can purchase after purchasing a base product, and consumers know the add-on price but mispredict their willingness to pay for it; and identify the limited ways in which heterogeneity in v affects our conclusions.

think of it as not participating in the market.

We look for Nash equilibria of the game played between firms, where—deviating from much of the literature—we impose that firms face a floor on the up-front price: $f_n \geq \underline{f}$. We assume that $\underline{f} \leq v$, so that consumers are willing to buy when they face an up-front price at the floor and believe that the additional price is zero. In stating our results, we focus on identifying conditions for and properties of deceptive equilibria—equilibria in which all firms shroud additional prices. Because no firm has an incentive to shroud if at least one firm unshrouds, there is always an unshrouded-prices equilibrium. When a deceptive equilibrium exists, however, it is more plausible than the unshrouded-prices equilibrium for a number of reasons. Most importantly, in Section 5.1 we show that in that case, the deceptive equilibrium is the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is. 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.

2.2 Motivation for and Interpretation of Key Assumptions

Our model has three key assumptions: that naive consumers might ignore the additional prices when making purchase decisions, that there is a floor on the up-front price, and that firms can costlessly and fully educate all consumers. We discuss these assumptions in turn, and provide examples of how our model fits various real-world markets.

The assumption that consumers might ignore some prices or fees is a simplified variant of many assumptions in behavioral industrial organization (DellaVigna and Malmendier 2004, Eliaz and Spiegler 2006, Grubb 2009, Heidhues and Kőszegi 2010, and others), and is consistent with observations in industries such as banking, credit-card, retail-investment, and mortgage services.[4] Furthermore, while we interpret the additional prices primarily as financial prices, our model applies

[4] For empirical evidence from a number of financial and other industries, see Hall (1997), Cruickshank (2000), Wilcox (2003), Barber, Odean and Zheng (2005), Agarwal, Driscoll, Gabaix and Laibson (2008), The Office of Fair Trading (2008), Stango and Zinman (2009), Federal Communication Commission (2010), and Woodward and Hall (2012). For some additional evidence from field experiments, see Hossain and Morgan (2006) and Chetty, Looney and Kroft (2009).

equally well to non-financial costs of owning a product that can be shrouded from consumers. For example, the product may be manufactured in disagreeable ways (e.g. in sweatshops or with environmentally unfriendly procedures), or it may be unhealthy or inconvenient to use.[5]

Our assumption of a price floor is based on a number of related arguments that have appeared in the literature. In Heidhues et al. (2011), we provide one microfoundation for the price floor based on the existence of "arbitrageurs" who would enter the market to make money off of a firm with overly low (for instance negative) prices, and who avoid the additional price because they are not interested in using the product itself. As a simple illustration in a specific case, consider the finding of Hackethal et al. (2010) that German bank revenues from security transactions amount to €2,560 per customer per year (2.43% of mean portfolio value). If a bank handed out such sums ex ante—even if it did so net of account maintenance costs—many individuals would sign up for (and then not use) bank accounts just to get the handouts. This threat creates a binding floor on banks' up-front price. In related models, Ko (2011) derives a version of our price floor from the presence of sophisticated consumers, Grubb (2012) and Armstrong and Vickers (2012) impose a no-negative-prices constraint in the context of a market with naive consumers, and Farrell and Klemperer (2007) discuss the same constraint in the context of switching-cost models.

Of course, in our model with a competitive market, we have imposed the price floor in an extreme form: there is a single number up to which firms are willing to lower the up-front price, but beyond which they cannot go. In reality, there is typically no such bright-line price floor. Generalizing from our setting, the intuitions for our main qualitative results on the role of wasteful and inferior products in generating profits require only that firms are less willing to cut the up-front price than a transparent total price, so that they make higher profits with shrouded than with unshrouded additional prices. This would be the case in any setting in which, similarly to our arbitrageurs model mentioned above and the model of add-on pricing by Ellison (2005), less

[5] In contrast to our model, in most of the above examples a consumer has some control over how much of the additional price she pays. So long as consumers' fundamental mistake is in underestimating additional prices, the logic of our model requires only that consumers cannot fully avoid these prices, so that firms can make profits on them. If consumers' mistake is in mispredicting their own behavior rather than prices, the model of Section 5.2 applies. Similarly, in many markets naive consumers may (contrary to our model) underestimate but not fully ignore additional prices. In this case, we can apply our model by thinking of the extra fees consumers expect to pay as being included in f_n, with the unexpected component of the extra fees being a_n.

profitable consumers are more responsive to the up-front price than more profitable consumers, and cutting a transparent total price mainly attracts more profitable consumers.

Importantly, although (as we will show) a deception-based positive-profit equilibrium may exist in our model when the floor on the up-front price binds, this does not mean that firms earn positive profits once their full economic environment is taken into account. Our stylized model focuses only on the stage of serving existing consumers, and ignores costs firm may have to pay to enter the industry, to advertise, to identify potential consumers, and so on. Nevertheless, since many industries motivating our analysis seem quite competitive even at the price-competition stage when entry costs have been sunk and potential consumers have been identified, the existence of positive profits at this stage is a potentially important message of our model.

Finally, our model assumes that firms can unshroud additional prices to consumers; in fact, one main goal of this paper is to identify market features that affect the incentive to educate consumers through unshrouding. In order to study these incentives in a theoretically clean and simple manner, we have imposed not only a competitive setting in production technologies (no firm is superior to all others in production cost), but also a kind of perfect competition in education: firms can costlessly reveal competitors' additional prices to all consumers. These assumptions are unrealistically extreme. In Section 5.1, we discuss various ways in which market power and costly unshrouding modify our results.

To put our model and its key assumptions into perspective, we briefly mention how to map it to several markets motivating our analysis. For banking services, the up-front price can correspond to initial charges or regular monthly fees and the additional price to overdraft fees and other contingent fees; for credit cards, the up-front price can correspond to the annual fee net of perks and the additional price to future interest payments and various fees and penalties; and for mutual funds, the up-front price can correspond to the front load and the additional price to future management fees. In each of these cases, the up-front price cannot drop much below zero without the danger of attracting unprofitable consumers. As a less clear-cut example, for mortgages the up-front price can correspond to initial monthly payments and the additional price to future monthly payments, prepayment penalties and other fees. In this case, it is unclear whether firms face a price floor.

But if—similarly to Ellison (2005)— cutting the initial monthly payments to very low levels would attract primarily risky borrowers, firms might not be willing to compete too much on this up-front price, imposing something akin to a price floor.

3 Profitable Deception

This section analyzes our basic model. We start in Section 3.1 with two benchmarks in which firms cannot earn positive profits by shrouding. In Section 3.2, we turn to our main result: we establish conditions under which an equilibrium with profitable deception can be maintained. In Section 3.3, we discuss the role of social wastefulness in facilitating deception.

3.1 Benchmarks: Non-Binding Price Floor or Sophisticated Consumers

First, we state what happens when the floor on the up-front price is not binding:[6]

Proposition 1 (Equilibrium with Non-Binding Price Floor). *Suppose $\underline{f} \leq c_{min} - \overline{a}$. Then, there exists a deceptive equilibrium. In any deceptive equilibrium, consumers buy the product from a most efficient firm, pay $f = c_{min} - \overline{a}, a = \overline{a}$, and get utility $v - c_{min}$, and firms earn zero profits.*

Since in a deceptive equilibrium consumers do not take into account additional prices when selecting 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 and behavioral-economics theories, firms compete aggressively for these valuable consumers ex ante, and bid down the up-front price until they eliminate net profits. In addition, since with these prices a firm cannot profitably undercut competitors and hence has no incentive to unshroud, a deceptive equilibrium exists. This equilibrium may be socially inefficient even though firms make zero profits: since consumers do not anticipate additional prices, they can be induced to buy a product whose value is below production cost (i.e., $v - c_{min}$ might be negative).

[6] Note that the proposition is stated in terms of what product consumers get rather than what firms do. Analogously to any standard Bertrand-competition model, there is an uninteresting multiplicity of equilibria due to the fact that a firm can make zero profits by charging the up-front price f identified in the proposition, as well as by charging a higher price and attracting no consumers. Equilibrium requires only that at least two firms charge the lowest price. Which of these equilibria obtains affects neither firm profits nor consumer welfare.

By Proposition 1, therefore, deception can occur in our model even if the price floor is not binding. Indeed, for some commonly invoked examples of deceptive products, such as hotel rooms and printers, the price floor does not seem to be binding. Nevertheless, for several reasons deception is likely to be less widespread and economically harmful in this case than in the case of a binding price floor below. First, while in the current case all consumers buy from a most-efficient firm and pay a total price equal to the firm's cost, in the case of a binding price floor they pay a higher price and also buy from less efficient firms. Second, the profits firms earn when the price floor is binding can potentially can potentially attract entry of less efficient producers, exacerbating productive inefficiency. Third, we show in our companion paper (Heidhues, Kőszegi and Murooka 2012) that the incentive to come up with deceptive products in the first place is very different with and without the price floor. With a non-binding price floor, firms have no incentive to invent new ways of charging additional prices, so they limit themselves to obvious deception opportunities. With a binding price floor, in contrast, the incentive to invent new hidden fees—even ones competitors can easily copy—is often very strong.

As a second benchmark, we consider equilibria when all consumers are sophisticated in that they observe the total prices $f_n + a_n$ and make purchase decisions based on these prices.

Proposition 2 (Equilibrium with Sophisticated Consumers)**.** *Suppose all consumers are sophisticated, and consider any $\underline{f}$. If $v > c_{min}$, then in any Nash equilibrium consumers buy the product at a total price of c_{min} from a most efficient firm. If $v < c_{min}$, then in any Nash equilibrium consumers do not buy the product. Firms earn zero profits in any equilibrium.*

Since sophisticated consumers understand the total price, firms cannot break even by selling a socially wasteful product to them. Furthermore, firms make zero profits in selling a socially valuable product as well: if there is a binding floor on the up-front price, firms simply switch to competing on the additional price, and—as there is no floor on this price—bid down the total price until profits are zero.

3.2 Naive Consumers with a Binding Price Floor

Taken together, Propositions 1 and 2 imply that for profitable deception to occur, both naive consumers must be present and the price floor must be binding. We turn to analyzing our model when this is the case, assuming for the rest of this section that all consumers are naive and $\underline{f} > c_n - \overline{a}$ for all n. To identify when a deceptive equilibrium exists, first note that if additional prices remain shrouded, 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 up-front price to $\underline{f}$. With consumers being indifferent between firms, firm n gets market share s_n and therefore earns a profit of $s_n(\underline{f} + \overline{a} - c_n)$. For this to be an equilibrium, no firm should want to unshroud additional prices. Once a firm unshrouds, consumers will be willing to pay exactly v for the product, so that firm n can make profits of at most $v - c_n$ by unshrouding and capturing the entire market. Hence, unshrouding is unprofitable for firm n if the following "Shrouding Condition" holds:

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

By extension, a deceptive equilibrium exists if (SC) holds for all n. Furthermore, since $s_n < 1$, Condition (SC) implies that $\underline{f} + \overline{a} > v$, so that in a deceptive equilibrium consumers receive negative utility. Proposition 3 states this result, and also says that if some firm violates Condition (SC), there is no deception in equilibrium:

Proposition 3 (Equilibrium with Binding Price Floor)**.** *Suppose $\underline{f} > c_n - \overline{a}$ for all n. If Inequality* (SC) *holds for all n, a deceptive equilibrium exists. In any deceptive equilibrium, $f_n = \underline{f}$ and $a_n = \overline{a}$ for all n, consumers receive negative utility, and firms earn positive profits. If Inequality* (SC) *is violated for some n, in equilibrium prices are unshrouded with probability one, consumers buy from a most efficient firm at a total price of c_{min}, and firms earn zero profits.*

The intuition for why firms might earn positive profits despite facing Bertrand-type price competition is in two parts. First, as in previous models and 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 better up-front terms. But the price floor prevents firms from competing away all profits from the additional price by lowering the up-front price.

Second, since firms cannot compete for consumers by cutting their up-front price, there is pressure for competition to shift to the additional price—but because competition in the additional price requires unshrouding, it is an imperfect substitute for competition in the up-front price. When a firm unshrouds and cuts the additional price by a little bit, consumers learn not only that the firm's product is the cheapest, but also that all products are more expensive than they thought. If $\underline{f} + \overline{a} > v$—that is, if consumers' utility from buying at the current total price is negative—this surprise leads consumers not to buy, so that the firm can attract consumers by unshrouding only if it cuts the additional price by a discrete margin. Since this may be unprofitable, the firm may prefer to shroud.

Importantly, if purchasing at the total market price is optimal (i.e., if $\underline{f} + \overline{a} \leq v$), then despite their surprise consumers are willing to buy from a firm that unshrouds and undercuts competitors' additional price by a little bit, so in this case a deceptive equilibrium does not exist. This logic indicates that deception in our model requires individual-welfare-reducing consumer purchases. And because firms can earn positive profits only through deception, our model also says that any profits must be associated with suboptimal consumer choices.[7]

Beyond showing that Condition (SC) is sufficient for profitable deception to occur, Proposition 3 establishes the (technically more difficult) converse that Condition (SC) is also necessary: if it is violated for some firm, then there is no deception in equilibrium, so that by classic Bertrand logic firms earn zero profits. The proof is by contradiction, and proceeds roughly as follows. If rivals shroud with positive probability, a firm can ensure positive profits by shrouding and choosing prices $\underline{f}, \overline{a}$. Since a firm that sets the highest total price when unshrouding has zero market share if some other firm unshrouds, to earn positive profits it must be that with positive probability all rivals set higher total prices when shrouding. For these high total prices, firms earn positive profits only when shrouding occurs, so that arguments akin to those above imply that they set prices $\underline{f}, \overline{a}$. But then a firm that violates Condition (SC) prefers to unshroud.

[7] Furthermore, if a deceptive equilibrium is played by firms, then productive efficiency also fails to hold: market shares are determined by how consumers happen to choose when indifferent. This contrasts sharply with natural specifications of classical Bertrand competition, where the market share of firms other than the most efficient ones is zero.

3.3 Socially Valuable versus Socially Wasteful Products

We now use Proposition 3 to identify some circumstances under which profitable deception does versus does not occur, distinguishing socially valuable and socially wasteful products.[8]

Non-vanishingly socially valuable product (there is an $\epsilon > 0$ such that $v > c_n + \epsilon$ 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), and thereby guarantees that only a zero-profit equilibrium exists. Intuitively, a firm that earns low profits from deception prefers to attract consumers through unshrouding, eliminating the possibility of profitable deception. This implies that an increase in the number of firms—which eventually leads some firm to violate Condition SC—can induce a regime shift from a high-price, deceptive equilibrium to a low-price, transparent equilibrium.[9]

Socially wasteful product ($v < c_n$ 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 $v < c_n$ for all n. Then, a profitable deceptive equilibrium exists.*

This perverse result has a simple and compelling logic: since a socially wasteful product cannot be profitably sold once consumers understand its total price, a firm can never profit from coming clean.

Some costly non-traditional mortgage products might be a good example for this case of our model. For instance, the Option Adjustable-Rate Mortgage allows borrowers to pay less than the interest for a period, leading to an increase in the amount owed and sharp (even 100-percent or

[8] We do not discuss in-between cases in which the product is valuable to produce by some firms but not other firms. The implications then depend on the market shares of efficient firms, and how these change with entry. For example, if there is an efficient firm whose market share approaches zero as the number of firms increases, the analysis of entry is akin to that in the case of a socially valuable product below.

[9] The reason for stating our result for *non-vanishingly* socially valuable products is that if the social value of a firm's product ($v - c_n$) 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 $s_n(\underline{f} + \overline{a} - c_n) < \epsilon < v - c_n$, in violation of Condition (SC).

higher) increases in monthly payments.[10] While this mortgage may make sense for consumers who confidently expect sharp increases in income or who are willing to take the risky gamble that house prices will appreciate, it likely served no purpose for many or most of the vast number of consumers who took it. Furthermore, by leading many consumers to overborrow and get into financial trouble, the product might well have lowered social welfare. Indeed, some features of Option ARMs, such as an introductory interest rate that applies for one or three *months*, serve only the purpose of deceiving borrowers about the product's cost. Our model says that Option ARMs continued to be sold and remained profitable in a seemingly competitive market not despite, but *exactly because* they were socially wasteful.

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. Second, our theory suggests a general competition-impairing feature in valuable industries that is not present in wasteful industries: to reduce the motive to deviate from their preferred positive-profit deceptive equilibrium in a valuable industry, each firm wants to make sure competitors earn sufficient profits from shrouding. This feature is likely to have many implications beyond the current paper (as, for instance, for innovation incentives in Heidhues et al. 2012), and implies that wasteful industries may sometimes be more fiercely competitive than valuable ones.

4 Sophisticated Consumers and Multi-Product Markets

Our analysis has so far assumed that all consumers are naive. In this section, we discuss the implications of assuming that some consumers are sophisticated in that they observe and take into account additional prices when making purchase decisions. We begin in Section 4.1 by pointing out how this change modifies the logic of our basic model, and then consider a multi-product market

[10] 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`. As one indication of how widespread Option ARMs had become, this product represented 19 percent of Countrywide's (the then-largest lender's) originations in 2005. The New York Times reports that Countrywide made gross profits of 4 percent on such loans, compared to profits of only 2 percent on traditional FHA loans (November 11, 2007).

in Section 4.2. Throughout this section, we assume that the proportion of sophisticated consumers is $\kappa \in (0,1)$, and that the price floor is binding: $\underline{f} > c_n - \overline{a}$ for all n.

4.1 Sophisticated Consumers in Our Basic Model

Notice that in any 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. Furthermore, note that if firm n unshrouds, it attracts consumers if and only if it cuts the total price to at most v—but if it does so, it attracts all naive and sophisticated consumers. Combining these considerations, unshrouding is unprofitable if

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

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

Condition (1) for the existence of a deceptive equilibrium has two notable implications. First, if the product is socially wasteful, the presence of sophisticated consumers does not affect our results, as a profitable deceptive equilibrium always exists. Intuitively, sophisticated consumers do not buy a socially wasteful product in equilibrium, so their presence is irrelevant—firms just attempt to exploit naive consumers. But second, if the product is socially valuable, the condition for a deceptive equilibrium to exist is stricter in the presence of sophisticated consumers. Intuitively, while these consumers do not buy the product when the additional price is high, they can be attracted by a price cut, creating pressure to cut the additional price—and by implication also to unshroud.

4.2 Sophisticated Consumers with an Alternative Transparent Product

We now move beyond Section 4.1 by assuming not only that there are sophisticated consumers, but also that there is another product in the market. Our analysis is motivated by the observation that in many markets, products that are more transparent than and seemingly superior to the deceptive products exist. Mutual-fund investors can choose low-cost index funds that will earn them higher returns than most managed funds. Many credit-card consumers could use debit cards for the same set of basic services and avoid most fees and interest. And many mortgage borrowers would be better served by simple traditional mortgages than by the complicated exotic products that have gained significant market share recently.

Formally, we modify our model above by assuming that each firm has an additional, transparent, product with value $w > 0$, where firm n's cost of producing product w is c_n^w. We let $\min_n\{c_n^w\} = c_{min}^w > 0$, and assume that there are at least two firms whose cost of producing product w is c_{min}^w. Crucially, we posit that product w is socially valuable ($w - c_{min}^w > 0$), and is not inferior to product v: $w - c_{min}^w \geq v - c_{min}$. Consumers are interested in buying at most one product. Firms simultaneously set the up-front and additional prices for product v, the single transparent price for product w, and decide whether to unshroud the additional price 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 and s_n^w, respectively. Then:

Proposition 4 (Profitability of Inferior Products). *Suppose $\underline{f} > c_n - \overline{a}$ for all n. For any shares s_n, s_n^w, there exists an equilibrium in which each firm shrouds the additional price of product v, naive consumers buy product v, and sophisticated consumers buy product w, if and only if $v - \underline{f} \geq w - c_{min}^w$. In such an equilibrium, firms sell the superior product to sophisticated consumers and earn zero profits on it, while they sell the inferior product to naive consumers and earn positive profits on it.*

Quite in contrast to the message of Section 4.1 that sophisticated consumers increase the pressure to unshroud, Proposition 4 says that if $v - \underline{f} \geq w - c_{min}^w$, a positive-profit equilibrium in which naive consumers are deceived always exists. This insight has a perverse implication: all it takes for profitable deception to occur is the availability of an *inferior* product that has a shroudable price component and a binding floor on the up-front price, and firms earn all their profits from selling the *inferior* product. The intuition for why the superior product guarantees a deceptive equilibrium with positive profits from the inferior product is in two parts. First, because sophisticated consumers realize that the deceptive product is costly but naive consumers believe it is a better deal, in equilibrium the two types of consumers separate. Second, if a firm unshrouded the additional price of the inferior product, consumers would realize that the other product is better, and would buy that product. As a result, a firm cannot make positive profits by unshrouding the additional price of the inferior product. In a sense, the superior product serves as a barrier to unshrouding the inferior product by rendering the inferior product socially wasteful in comparison.

The condition $v - \underline{f} \geq w - c^w_{min}$ for a positive-profit deceptive equilibrium to exist is a sorting condition: it implies that because they ignore its additional price, naive consumers mistakenly find the inferior product v more attractive than the superior product w. This condition holds if product w is not much better than product v or $\underline{f}$ is not too high. For instance, although a naive consumer may realize that a debit card fulfills the same functions that she uses in a credit card, she may still prefer a credit card because she falsely believes that its perks (e.g. cash-back bonuses) make it a better deal.[11]

As an example, consider the market for mutual funds. Although not everyone agrees with this view, many researchers believe that because few mutual-fund managers can persistently outperform the market by enough to make up for their high fees (Carhart 1997, Kosowski, Timmermann, Wermers and White 2006), most managed funds are inferior to index funds. As a result, the explosion of managed funds is often seen as a puzzle (Gruber 1996, French 2008, for example). Our model says that managed funds could have remained profitable (and hence have attracted a lot of entry) not despite, but *exactly because* index funds that are superior to them exist.

The conclusions of Proposition 4 continue to hold if we assume that consumers misperceive the inferior product's value rather than its price. Suppose that product v has no additional price, but consumers have false beliefs about its value: they believe the value is v, but it is actually $v - \overline{a}$. Continuing with the mutual-fund example, naive investors might overestimate the ability of a manager to pick good investments rather than underestimate the fees she charges. Even then, if

[11] Although we have exogenously imposed that product w is transparent, this will often arise endogenously even if firms make an unshrouding decision regarding both products. Clearly, under the condition of Proposition 4, an equilibrium in which product v is shrouded and product w is unshrouded exists in that case as well. If in addition $w > v$ and there are sufficiently many firms in the market, the *only* profitable equilibrium is the one in which the superior product is unshrouded and the inferior product is shrouded. Consider, for example, a candidate equilibrium in which the superior product w is shrouded. Then, naive consumers must be buying product w; otherwise, a firm could attract all these naive consumers by setting prices $\underline{f}, \overline{a}$ on product w, and for a low-profit firm this would be a profitable deviation. But if naive consumers are buying product w, a low-profit firm has an incentive to unshroud product w in order to capture this socially valuable market.

product w is superior—that is, if $w - c^w_{min} \geq v - \overline{a} - c_{min}$—Proposition 4 and the logic behind it survive unchanged.

Note also that Proposition 4 holds for *any* market shares s_n, s^w_n for the two products. In particular, this means that in our model not even a "specialist" in the superior product—a firm that sells exclusively or 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. In Section 5.1, we discuss how this result is qualified with market power and costly unshrouding.

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, firms may pay intermediaries to convince consumers to buy the inferior product.[12] Second, 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 to an inferior good. Third, firms may inform consumers unaware of the inferior product of the product's existence, yet not do the same for the superior product. Fourth, firms may make costly (real or perceived) improvements to the inferior product to make it more attractive to consumers.

5 Extensions and Modifications

To demonstrate some robustness of our findings, as well as to raise additional issues, in this section we discuss various extensions and modifications of our framework. Unless otherwise stated, we continue to assume that the price floor is binding.

[12] Indeed, Anagol, Cole and Sarkar (2011) and Mullainathan, Nöth and Schoar (2011) document that intermediaries tend to disproportionately push inferior products in the life-insurance and mutual-fund markets, respectively, and do so because they receive higher commissions from firms. Consistent with the perspective that this hurts consumers, Bergstresser, Chalmers and Tufano (2009) find that broker-sold funds deliver lower risk-adjusted returns than do direct-sold funds.

5.1 Costly Unshrouding and Market Power

So far, we have assumed a competitive market in both production and unshrouding technologies: no firm has strictly lower cost than all others, and firms can unshroud costlessly. In this section, we discuss implications of relaxing these assumptions.

Costly unshrouding. Consider the same game as in Section 2, except that firm n has to pay $\eta \geq 0$ to unshroud additional prices, and such unshrouding educates only a fraction λ_n of naive consumers. We begin with the special case $\lambda_n = 1$, and provide one justification for our presumption that firms play a deceptive equilibrium whenever it exists:

Proposition 5 (Unique Equilibrium with Costly Unshrouding)**.** *Fix all parameters other than η, and suppose that $\underline{f} > c_n - \overline{a}, \lambda_n = 1$ 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.*

Proposition 5 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. To see the logic of this result, notice first that if $\eta > 0$, in order to unshroud a firm must make positive gross profits afterwards. Hence, no firm unshrouds with probability one—as this would lead to Bertrand-type competition and zero gross profits. Now for each firm, 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.

To draw out further implications of costly unshrouding, we consider general $\lambda_n \leq 1$, allowing λ_n to differ across firms to capture the notion that some firms (e.g. firms with established marketing departments) may at the same cost be able to reach more consumers than other firms. We assume that if firms $n_1, \ldots, n_k$ unshroud, a fraction $\max\{\lambda_{n_1}, \ldots, \lambda_{n_k}\}$ of consumers becomes educated. In addition, we suppose that a firm can offer only one pricing scheme f, a.[13] Then, the most tempting deviation from a candidate deceptive equilibrium with prices $\underline{f}, \overline{a}$ is to unshroud, still charge $f_n = \underline{f}$ to attract a share s_n of the remaining uneducated consumers, and lower a_n to $v - \underline{f}$ to attract all educated consumers. This is unprofitable if the following variant of the Shrouding Condition (SC) holds:

$$s_n(\underline{f} + \overline{a} - c_n) \geq [\lambda_n + (1 - \lambda_n)s_n](v - c_n) - \eta. \tag{2}$$

Costly unshrouding somewhat qualifies the distinction between socially wasteful and socially valuable products we have emphasized in Section 3.3. By Condition (2), it is still true that if the product is socially wasteful, a deceptive equilibrium exists for any number of firms. It is no longer true, however, that with a socially valuable product and sufficiently many firms, a deceptive equilibrium does not exist: if unshrouding is expensive or very partial, no firm may prefer to do it even if it can sell the product at positive gross profits. Nevertheless, a version of our basic message continues to hold: the worse a product is, the more likely it is that a deceptive equilibrium exists.

While we have not characterized equilibria in general when Condition (2) is violated for some firm, an interesting equilibrium arises in the special case in which this condition is violated for some firm n but

$$(1 - \lambda_n)\frac{s_{n'}}{1 - s_n}(\underline{f} + \overline{a} - c_{n'}) \geq \max\{[\lambda_{n'} + (1 - \lambda_{n'})s_{n'}](v - c_{n'}) - \eta, [\lambda_n + (1 - \lambda_n)s_{n'}](v - c_{n'})\} \tag{3}$$

for all $n' \neq n$. Then, there is an equilibrium in which firm n educates a fraction λ_n of consumers, while (by Condition (3)) all other firms sell to naive consumers. This equilibrium seems consistent with the observation that some firms have attempted to unshroud and sell transparent products to consumers, but they have often had only limited impact on their market. Our theory makes comparative-statics predictions on which firm is most likely to use a transparent strategy. Rewriting the reverse of Condition (2) as $\lambda_n(1 - s_n)(v - c_n) - s_n(\underline{f} + \overline{a} - v) - \eta > 0$, we obtain that (ceteris paribus) a firm is more likely to offer a transparent product if (i) it can educate more consumers

[13] If a firm can offer multiple pricing schemes, from the perspective of deriving conditions for a deceptive equilibrium to exist we can think of educated and uneducated consumers as being in separate markets. Because deviating from deceptive pricing in the market for uneducatable consumers is obviously unprofitable, a firm's incentives to change its prices and to unshroud derive entirely from the market for educatable consumers. Hence, this case is equivalent to that above for $\lambda_n = 1$. Incorporating the unshrouding cost η, the condition for an equilibrium in which all firms always shroud to exist becomes $s_n(\underline{f} + \overline{a} - c_n) \geq (v - c_n) - \eta$ for any n.

(λ_n is high); (ii) it, such as a new entrant, has a smaller market share (s_n is low); or (iii) it is more efficient (c_n is lower).[14]

While it qualifies our results on single-product valuable industries, the addition of an unshrouding cost does not affect the main message of Section 4 that there is often an equilibrium in which naive consumers buy a deceptive inferior product: since a firm has no incentive to unshroud the inferior product even if this is free and full, it certainly does not have an incentive to do so if the same thing is costly and partial.

Market power. We now consider the effect of market power of a very simple form—we assume that firm n has strictly lower cost than any other firm: $c_n < \min_{n' \neq n} c_{n'}$. We also suppose first that $\eta = 0$. In this case, the condition for when a deceptive equilibrium exists in our basic model remains unchanged. Intuitively, since when prices are shrouded competitors' total prices are not sensitive to their costs, these costs do not play a role in determining whether firm n wants to unshroud.

Market power does have an interesting effect on outcomes in our multiple-products model. In particular, suppose that firm n has market power in the market for the superior product ($c^w_{min} - c^w_n \equiv M > 0$ where $c^w_{min} \equiv \min_{n' \neq n} c^w_{n'}$). For this market, we make the common assumption that no firm charges a price below cost, so that in equilibrium firm n charges c^w_{min} and attracts all consumers. Then:

Proposition 6 (Market Power in the Superior Product)**.** *Suppose $\underline{f} > c_{n'} - \overline{a}$ for all n', $\eta = 0$, and $v - \underline{f} \geq w - c^w_n$.*

I. If $M > s_n(\underline{f} + \overline{a} - c_n)$, there is an equilibrium in which firm n unshrouds the additional price of product v and educates a fraction λ_n of naive consumers, other firms do not unshroud, sophisticated consumers and consumers educated by firm n buy product w, and uneducated naive consumers buy product v.

II. If $M \leq s_n(\underline{f} + \overline{a} - c_n)$, there exists an equilibrium in which each firm shrouds the additional price of product v, sophisticated consumers buy the transparent product w, and naive consumers buy the shrouded product v.

[14] If Condition (2) is violated for multiple firms, a free-riding issue arises: a firm prefers another firm to pay the unshrouding cost, even if it then prefers to compete in price for educated consumers. This implies that no pure-strategy equilibrium exists, and it is difficult to characterize the equilibrium in general.

Part I of Proposition 6 says that if firm n's market power is larger than its share of deceptive profits, it prefers to unshroud and attract the naive consumers it can educate. Such a situation could arise, for instance, if firm n sells only the superior product ($s_n = 0$). This observation explains why some specialists in superior products—e.g., Vanguard in the market for mutual funds—have tried to educate consumers about the inferiority of alternative products. At the same time, Part II of Proposition 6 says that if firm n has a large share in the market for the inferior product or the margin on this product is large ($s_n(\underline{f} + \overline{a} - c_n)$ is high), it prefers not to educate consumers even if it has market power in the superior product. Similarly to Section 4.2, the intuition is that deception creates a large profit margin in the inferior product despite competition, and unshrouding the additional price of this product only leads consumers to buy the less profitable superior one.

Market power and costly unshrouding. Going further, the extent of unshrouding can be even more limited if unshrouding is costly ($\eta > 0$). Suppose, as often seems to be the case in reality, that competition in the superior product is relatively fierce (M is small). Then, firm n is unwilling to educate naive consumers unless η is also small. And going slightly beyond our model, 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 limited education. These considerations may help explain why consumer education is often limited.

5.2 Misprediction of Add-On Demand

As an alternative to our specification of consumer naivete above, in this section we analyze a model in which a consumer underestimates not the total price of the product, but her own demand for some add-on to the product. When getting a credit card, for example, a consumer may be aware that she will face a high interest rate on any long-term debt she carries, but incorrectly expect to pay off her outstanding debt within a short period.[15]

We use the same model as in Section 2, with the following modifications. Instead of assuming that f_n and a_n are two components of a product's price, we posit that f_n is the price of a base

[15] Consistent with this example, Ausubel (1991) finds that consumers are much less responsive to the post-introductory interest rate in credit-card solicitations than to the teaser rate, even though the former is more important in determining the amount of interest they will pay.

product (e.g., the convenience use of a credit card) and a_n is the price of an add-on (e.g., long-term borrowing on the credit card). A consumer can only buy a firm's add-on if she purchased that firm's base product. We assume that consumers know a_n, but have false beliefs about their demand for the add-on: whereas their actual willingness to pay will be $\overline{a}$, they believe their willingness to pay will be $\hat{a} < \overline{a}$. Consumers value the product with the add-on at v; hence, their perceived value for the product without the add-on is $v - \hat{a}$.[16] In contrast to our assumption in Section 2 that any firm can eliminate consumer misperceptions, in this version of the model we do not assume that firms can do so. This reflects our view that convincingly explaining to a consumer how she herself will behave is more difficult than highlighting a price; indeed, a consumer may be presented with and readily believe information about how the average consumer behaves, but still think that this does not apply to her.

Proposition 7 identifies the key result in this variant of our model. As in the rest of the paper, we identify conditions for profitable equilibria in which consumers mispredict how much they will pay. But because consumers understand the add-on price when they buy the base product, we refer to such an equilibrium as an "exploitative equilibrium" rather than a deceptive equilibrium.

Proposition 7 (Equilibrium in Underestimation-of-Demand Model)**.** *Suppose $\underline{f} > c_n - \overline{a}$ for all n. In any exploitative equilibrium, $f_n = \underline{f}$ and $a_n = \overline{a}$ for all n. An exploitative equilibrium exists if and only if*

$$s_n(\underline{f} + \overline{a} - c_n) \geq \underline{f} + \hat{a} - c_n \text{ for all } n \in \{1, \ldots, N\}. \tag{4}$$

The underestimation-of-demand model shares the prediction of our basic model above that sometimes a profitable exploitative equilibrium exists despite price competition in undifferentiated products, and this equilibrium relies on consumers misunderstanding what they are buying. But

[16] An alternative way to set up the model is to assume that the consumer's value for the product without the add-on is v, with her perceived willingness to pay for the add-on still being $\hat{a}$. The two formulations generate the same predictions, but have slightly different interpretations. In the former case, the consumer overestimates her value for the product without the add-on. For instance, a mobile-phone consumer might not realize how painful it is to forego calling while traveling in areas where roaming charges apply. In the latter case, the consumer understands the value of the product without the add-on, but does not realize how tempted she will be buy the add-on. For example, a consumer may understand the convenience value of a credit card, but underappreciate her tendency to borrow on it.

the mechanism is somewhat different. Because consumers do not believe they will buy the add-on at a price above $\hat{a}$, they do not respond to a firm that undercuts competitors' add-on price of $\overline{a}$ by a little bit. Instead, to attract consumers a firm must cut its add-on price discretely to $\hat{a}$—the price at which consumers *believe* they will want the add-on—and this may not be worth it. Condition (4) for an exploitative equilibrium to exist says that firm n makes more profits charging the highest add-on price $\overline{a}$ and getting market share s_n than charging only the add-on price $\hat{a}$ and getting all consumers.[17]

Given the similarity of Conditions (SC) and (4), the implications of Proposition 7—as well as those of introducing sophisticated consumers—are also similar to those of the basic model. These implications, however, now depend not on whether the product is socially wasteful, but on whether the product is unprofitable to sell when charging the add-on price at which consumers *think* they will value the add-on (i.e. whether $\underline{f} + \hat{a} < c_n$). If the product is profitable to sell at this "virtual" price, then with a sufficient number of firms at least one is willing to lower the add-on price to the virtual price, eliminating the exploitative equilibrium. But if the product is unprofitable to sell at the virtual price, then a profitable exploitative equilibrium exists independently of the number of firms in the industry or other parameter values.

An example consistent with the above prediction on when entry does not eliminate profitable exploitative practices may be the credit-card market. Suppose, for instance, that consumers ignore the 18% interest rate on credit-card balances because they believe they will not carry a balance for interest rates exceeding 5%. Then, to attract consumers a firm must cut its interest rate to 5%, and this may be unprofitable.

[17] The above equilibrium is not robust to assuming that consumers perceive the probability of consuming the add-on to be positive, no matter how small the probability is. With products being perfect substitutes, consumers then respond to any decrease in the add-on price, so an equilibrium with an add-on price of $\overline{a}$ does not exist. Even so, if there is a positive measure of consumers who perceive the probability of purchasing the add-on at a price of $\overline{a}$ to be zero, a positive-profit mixed-strategy exploitative equilibrium exists because—similarly to the "captive" consumers in Shilony (1977) and Varian (1980)—these consumers provide a profit base that puts a lower bound on firms' total profits. Furthermore, it is clear that these profits can be sufficient to deter unshrouding.

5.3 Further Extensions and Modifications

For simplicity, our basic model assumes that all consumers have the same valuation v for the product, but our main insights survive when there is heterogeneity in v. As an analogue of Proposition 3, a deceptive equilibrium with prices $\underline{f}, \overline{a}$ often exists because unshrouding would lead consumers with values between $\underline{f}$ and $\underline{f} + \overline{a}$ not to buy, discretely reducing industry demand. The deceptive equilibrium is more likely to exist when there are more such consumers—that is, when there are more consumers who are mistakenly buying the product. And a deceptive equilibrium exists whenever the product could not be profitably sold to consumers who understand its total price. As above, this is the case whenever the product is socially wasteful to produce, for example because no consumer values it above marginal cost, or (in a natural extension of our model) the number of such consumers is insufficient given some fixed costs of production. But if the product can be profitably sold in a transparent way, then with a sufficient number of firms at least one firm would choose to unshroud, eliminating the deceptive equilibrium.

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.[18]

[18] The main results of our paper are also robust to allowing the maximum additional price to be different across firms, with firm n being able to set $\overline{a}_n$; in fact, our proof of Propositions 3 and 5 in the appendix allows for this possibility. This assumption substantively modifies only Proposition 1 on equilibria with a non-binding price floor: because firms that are better at exploiting consumers can afford lower up-front prices, it is now not the firms with the lowest c_n that sell to consumers in a deceptive equilibrium, but the firms with the lowest $c_n - \overline{a}_n$. This creates allocative inefficiency even though firms make zero profits.

Finally, while in our basic analysis we assume that the price floor is binding ($\underline{f} > c_n - \overline{a}$) for *all* firms, the same qualitative points regarding the existence and properties of a positive-profit deceptive equilibrium survive if the price floor is binding for at least two, but not necessarily all, firms. Consider a version of Condition (SC) in which shares are adjusted assuming that only firms with $\underline{f} > c_n - \overline{a}$ are in the market. If this modified condition holds for all firms with $\underline{f} > c_n - \overline{a}$, then there is a positive-profit deceptive equilibrium in which the other firms do not sell: since

6 Related Theoretical Literature

In this section, we discuss theories most closely related to our paper. Relative to this literature, our main contribution is to identify the central role of socially wasteful and inferior products in maintaining profitable deception.

In Gabaix and Laibson's (2006) model, consumers buy a base good and can then purchase an add-on whose price might be shrouded by firms, and can also take costly steps to avoid the add-on. Gabaix and Laibson's main prediction is that unshrouding can be unattractive because it turns profitable naive consumers (who buy 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 draw out a number of new implications.[19]

In research complementary to ours, Grubb (2012) considers services, such as a mobile phone or a bank account with overdraft protection, for which consumers may not know the marginal price, 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.[20]

Our result that firms sell profitable inferior products to unknowing consumers may seem rem-

Condition (SC) can only be satisfied if $\underline{f} + \overline{a} > v$, which together with $\underline{f} \leq c_n - \overline{a}$ implies that $c_n > v$, a firm that is out of the market cannot profitably unshroud and undercut competitors.

[19] Relatedly, Piccione and Spiegler (2010) 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 (2011) and Kosfeld and Schüwer (2011) demonstrate that educating naive consumers in a Gabaix-Laibson framework can decrease welfare because formerly naive consumers may engage in inefficient substitution of the add-on.

[20] Our theory also builds on a growing literature in behavioral industrial organization that assumes consumers are not fully attentive, mispredict some aspects of products, or do not fully understand their own behavior. 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).

iniscent of a similar potential implication in classical asymmetric-information models. But the mechanism in our model based on the incentive to unshroud is different, and hence generates different comparative statics. For instance, while in a rational asymmetric-information setting a lower-quality inferior product may be more profitable to sell than a superior product because it is cheaper to produce and consumers do not know its value, our theory predicts that a lower-quality inferior product may be more profitable than a superior product even if it is more expensive to produce and consumers do know its value. In assuming that consumers can be induced to pay high additional fees once they buy a product, our theory also 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 natural specifications of a rational switching-cost model. Most importantly, if consumers know or learn product attributes, a rational switching-cost model predicts that a firm is better off selling a superior rather than an inferior product, so that such a model does not predict the systematic sale of inferior products in competitive markets.[21] In addition, this type of model does not predict a consideration analogous to the threat of unshrouding by competitors, and hence it does not generate our results regarding the effect of entry.[22]

7 Some Policy Implications and Conclusion

While the main goal of this paper is to explore features of markets for deceptive products, our insights have some immediate policy implications—and call for exploring more of these implications.

[21] 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, 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.

[22] Similarly, when consumers must pay classical search costs to find out prices (or product features), at a broad level one can think of the prices as being partly shrouded. While we believe that search costs are extremely important in the markets we consider, by themselves they do not seem to fully explain why inferior or socially wasteful products are systematically sold in a more profitable way than superior products. Furthermore, although we have no precise empirical evidence, it does not seem that firms are playing the mixed-strategy pricing equilibrium predicted by these models.

As an important example, consider the impact of a policy that decreases the maximum additional price firms can charge from $\overline{a}$ to $\overline{a}' < \overline{a}$ in the range where the price floor is binding.[23] 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. 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. In addition, a decrease in the additional price can lead to some firm violating Condition (SC), in which case the market becomes transparent, prices drop further, and productive efficiency obtains.

Relatedly, the regime shift from deceptive to transparent pricing predicted by our model for socially valuable products identifies a consumer-protection benefit of competition policies that increase the number of firms in the market. Nevertheless, because 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 in our model.

An important agenda for future research is analyzing the potential impact of costly education campaigns by a social planner or consumer group, especially how a consumer group would finance such a campaign and whether it could compete with firms. If consumers can solve the free-rider problem and organize a consumer group to educate, presumably firms can also solve their own free-rider problem and organize an interest group to obfuscate—and the latter group will have more money behind it. With multiple institutions attempting to provide conflicting advice, naive consumers may find it difficult to sort out whom they should believe.

Finally, in this paper we have 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

[23] This decrease could come from regulation that restricts the extent to which firms can overcharge consumers ex post. Although such regulation seems extremely hard in practice due to the difficulty of precisely defining what overcharging means, it may be possible in specific cases. For example, the Credit Card Accountability, Responsibility, and Disclosure (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 up-front price—akin to recent regulations of European low-cost airlines—can also serve to decrease $\overline{a}$.

can be thought of as a form of innovation. In our companion paper (Heidhues et al. 2012), we study the incentives for such "exploitative innovation," and contrast them with innovation that benefits consumers. Here too we find a perverse incentive: because learning ways to charge consumers higher hidden fees 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. In contrast, the incentive to make an innovation that increases the product's value to consumers is zero or negative if competitors can copy the innovation, and even if they cannot the incentive is strong only when the product is socially wasteful. The possibility of exploitative innovation also adds caution to our conclusion that policy should aim to lower the maximum additional price: such policy can greatly increase firms' incentive to make new exploitative innovations, and hence may have a small net effect.

References

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson, "Learning in the Credit Card Market," 2008. NBER Working Paper #13822.

Anagol, Santosh, Shawn Cole, and Shayak Sarkar, "Understanding the Incentives of Commissions Motivated Agents: Theory and Evidence from the Indian Life Insurance Market," 2011. Mimeo, Wharton.

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

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

Barber, Brad M., Terrance Odean, and Lu Zheng, "Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows," *Journal of Business*, 2005, *78* (6), 2095–2119.

Bergstresser, Daniel, John M. R. Chalmers, and Peter Tufano, "Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry," *Review of Financial Studies*, 2009, *22* (10), 4129–4156.

Carhart, Mark M., "On Persistence in Mutual Fund Performance," *Journal of Finance*, 1997, *52* (1), 57–82.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Cruickshank, Don, "Competition in UK Banking: A Report to the Chancellor of the Exchequer," Technical Report, HM Treasury 2000.

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

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

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

___ **and Sara Fisher Ellison**, "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 2009, *77* (2), 427–452.

Farrell, Joseph and Paul Klemperer, "Coordination and Lock-In: Competition with Switching Costs and Network Effects," in Mark Armstrong and Rob Porter, eds., *Handbook of Industrial Organization*, Vol. 3, North-Holland, 2007.

Federal Communication Commission, "Comment Sought on Measures Designed to Assist U.S. Wireless Consumers to Avoid "Bill Shock"," Technical Report 2010. Public Notice May 11 2010. CG Docket No. 09-158.

French, Kenneth R., "Presidential Address: The Cost of Active Investing," *Journal of Finance*, 2008, *63* (4), 1537–1573.

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

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

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

Gruber, Martin J., "Another Puzzle: The Growth in Actively Managed Mutual Funds," *Journal of Finance*, 1996, *51* (3), 783–810.

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

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

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

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

——, ——, **and** ——, "Exploitative Innovation," 2012. Working Paper in Preparation.

Hossain, Tanjim and John Morgan, "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *B.E. Journal of Economic Analysis and Policy: Advances in Economic Analysis and Policy*, 2006, *6* (2), 1–27.

Ko, K. Jeremy, "Disclosure and Price Regulation in a Market with Potentially Shrouded Costs," 2011. Mimeo.

Kosfeld, Michael and Ulrich Schüwer, "Add-on Pricing, Naive Consumers, and the Hidden Welfare Costs of Education," 2011. CEPR Working Paper No. 8636.

Kosowski, Robert, Allan Timmermann, Russ Wermers, and Hal White, "Can Mutual Fund "Stars" Really Pick Stocks? New Evidence from a Bootstrap Analysis," *Journal of Finance*, 2006, *61* (6), 2551–2595.

Laibson, David I. and Leeat Yariv, "Safety in Markets: An Impossibility Theorem for Dutch Books," 2007. Working Paper, Caltech.

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

Mullainathan, Sendhil, Markus Nöth, and Antoinette Schoar, "The Market for Financial Advice: An Audit Study," 2011. Mimeo, Massachusetts Institute of Technology.

Office of Fair Trading, "Personal Current Accounts in the UK," Technical Report 2008.

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

Shilony, Yuval, "Mixed Pricing in Oligopoly," *Journal of Economic Theory*, 1977, *14*, 373–388.

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

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

Stango, Victor and Jonathan Zinman, "What Do Consumers Really Pay on Their Checking and Credit Card Accounts? Explicit, Implicit, and Avoidable Costs," *American Economic Review Papers and Proceedings*, 2009, *99* (2), 424–429.

Varian, Hal R., "A Model of Sales," *American Economic Review*, 1980, *70* (4), 651–659.

Wilcox, Ronald T., "Bargain Hunting or Star Gazing? Investors' Preferences for Stock Mutual Funds," *The Journal of Business*, 2003, *76* (4), 645–663.

Woodward, Susan E. and Robert E. Hall, "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence," *American Economic Review*, 2012, *100* (2), 511–515.

Appendix: Proofs

Proof of Proposition 1. First, we establish properties of any deceptive equilibrium. Note that in any such equilibrium, for almost any f_n for which firm n has positive expected market share, $a_n = \overline{a}$. Since the additional price of a firm that has zero market share does not affect the behavior and outcomes of consumers or any other firm, when solving for the prices at which consumers buy and profits we can assume without loss of generality that $a_n = \overline{a}$ also for other prices. Now consider a model of Bertrand competition in which firm n has cost $c_n - \overline{a}$ and is choosing f_n. For any strategy profile, this game generates the same profits as in a modification of our game in which firms are restricted to shrouding. Hence, the standard Bertrand proof implies that consumers buy the product from a most efficient firm and pay $a = \overline{a}, f = c_{min} - \overline{a}$; ex-post utility of consumers is $v - c_{min}$; and firms earn zero profits.

Now each firm setting $f_n = c_n - \overline{a}$ and shrouding is clearly an equilibrium. □

Proof of Proposition 2. With consumers who observe and take into account the additional price, we have Bertrand competition in the total price. □

Proof of Proposition 3. We establish a slightly more general version of this proposition: we allow the maximum additional prices firms can impose to differ across firms. Let $\overline{a}_n$ be the maximum additional price firm n can impose. We prove the statement of Proposition 3 with Inequality (SC) replaced by

$$s_n(\underline{f} + \overline{a}_n - c_n) \geq v - c_n. \tag{5}$$

We have argued in the text that in any deceptive equilibrium firm n sets its maximal additional price, which now is $\overline{a}_n$. The same argument as in the text also establishes that if Inequality (6) holds for all n, then there is a deceptive equilibrium in which all firms set $\underline{f}, \overline{a}_n$. We now provide a formal argument for why firms set $\underline{f}$ 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 set the additional price $\overline{a}_n$. Note that by setting $f_n = \underline{f}$, firm n can guarantee itself a profit of $s_n(\underline{f} + \overline{a}_n - c_n) > 0$. As a result, no firm will set a price $f_n > v$, because then no consumer would buy from it. Take the supremum $\overline{f}$ of the union of the supports of firms' up-front price distributions. We consider two cases. First, suppose that some firm sets $\overline{f}$ with positive probability. Then, all firms have to set $\overline{f}$ with positive probability; otherwise, a firm setting $\overline{f}$ would have zero market share with probability one. Then, we must have $\overline{f} = \underline{f}$; otherwise, a firm could profitably deviate by moving the probability mass to a slightly lower price. Second, suppose that no firm sets $\overline{f}$ with positive probability. Suppose firm n's price distribution has supremum $\overline{f}$. Then, as f_n approaches $\overline{f}$, firm n's expected market share and hence expected profit approaches zero—a contradiction.

In the following, we prove by contradiction that if Inequality (6) is violated for some firm, then in any equilibrium additional prices are unshrouded with probability one. Note that if unshrouding occurs with probability one, then we have Bertrand competition in the total price, and hence consumers buy from a most efficient firm at a total price of c_{min} and all firms earn zero profits. The proof that unshrouding occurs with probability one proceeds in three steps.

(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, a firm can guarantee itself positive profits by shrouding and offering the contract $\underline{f}, \overline{a}_n$, which attracts consumers since $v > \underline{f}$ and makes positive profits since $\underline{f} + \overline{a}_n > c_n$ for all n.

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

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

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

If only one firm m sets $\hat{f}'$ with positive probability, then firm m has zero market share both when unshrouding occurs or when shrouding occurs and some firm other than m sets a total price strictly greater than $\hat{t}$. Because firm m earns positive profits by (i) and is the only firm that sets $\hat{f}'$ with positive probability conditional on the total price being strictly higher than $\hat{t}$, every firm except for m shrouds and sets its up-front fee strictly higher than $\hat{f}'$ and its total price weakly lower than $\hat{t}$ with positive probability. Suppose first $m = n$. Then, there exists a firm $l \neq n$ that shrouds and sets an up-front fee $f_l > \hat{f}'$, $a_l \leq \hat{t} - f_l$ with positive probability. Since $\overline{a}_l > \hat{t} - \hat{f}'$, firm l can increase its profits by decreasing all prices $f_l > \hat{f}'$ to $\hat{f}'$ and increasing its additional price holding the total price constant—a contradiction. Next, suppose $m \neq n$. Then, firm n shrouds and sets $f_n > \hat{f}'$ with positive probability and charges an additional price $a_n \leq \hat{t} - f_n$ with probability 1 when charging these up-front prices because $\hat{f}'$ is the supremum of the up-front price conditional on charging a total price strictly above $\hat{t}$. For almost all of these up-font prices, firm n must earn strictly positive profits when shrouding occurs; otherwise firm n could increase its profits by unshrouding prices for which it earns no profits when shrouding occurs and guarantee itself positive profits also when all rivals $n' \neq n$ shroud and charge a total price above $\hat{t}$. Thus, firm m shrouds and sets $f_m \geq f_n > \hat{f}'$, $a_m \leq \hat{t} - f_m$ with positive probability. Since $\overline{a}_m > \hat{t} - \hat{f}'$, firm m can increase its profits by decreasing all prices $f_m > \hat{f}'$ to $\hat{f}'$ and increasing its additional price holding the total price constant—a contradiction.

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

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

Now suppose that firm n charges an up-front 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 or equal to $\hat{t}$ when shrouding. But if it unshrouds and sets the same prices, it would also earn profits when all rivals shroud and set a price above $\hat{t}$, thereby strictly increasing its profits—a contradiction. Hence firm n also must set $\underline{f}$ with probability one when shrouding.

Second, suppose not all firms other than n set a total price strictly above $\hat{t}$. Hence, 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' up-front prices $\hat{f}'$ conditional on the total price being greater than or equal to $\hat{t}$. The remainder of the proof is the same as above.

(iii): *Additional prices are unshrouded with probability one.* Suppose not. Then, each firm chooses to shroud with positive probability. Take the infimum of total prices $\tilde{t}$ set by any firm when shrouding. We consider two cases. First, suppose $\tilde{t} \leq v$. Take a firm that achieves the infimum. By (i), this firm earns positive profits. For any $\epsilon > 0$, take total prices below $\tilde{t} + \epsilon$ of the firm. By unshrouding and setting $\tilde{t} - \epsilon$, the firm decreases its profits by at most 2ϵ when one or more other firms unshroud, but discretely increases its market share if all other firms shroud. Hence, for a sufficiently small $\epsilon > 0$ this is a profitable deviation—a contradiction. Second, suppose $\tilde{t} > v$. Take firm n that violates Inequality (6). By (ii), firm n charges the up-front price $\underline{f}$ whenever it shrouds. Note that firm n's profits are zero when a rival decides to unshroud, and its profits are at most $s_n(\underline{f} + \overline{a}_n - c_n)$ when shrouding occurs. But then, deviating and setting a total price to v is profitable by Inequality (6), because conditional on others shrouding it would earn $v - c_n$. $\square$

Proof of Proposition 4. First, we show that if an equilibrium of the type identified in the proposition exists, then $v - \underline{f} \geq w - c^w_{min}$. Since in such an equilibrium sophisticated consumers are buying the transparent product, standard Bertrand-competition logic implies that the total price of product w is c^w_{min} and firms earn zero profits on w. For product v, in turn, the same argument as in Proposition 3 shows that in a deceptive equilibrium in which naive consumers buy this product firms choose the up-front price $\underline{f}$ and additional price $\overline{a}$. Then, naive consumers' ex-ante perceived utility from buying v is $v - \underline{f}$ and their ex-ante perceived utility of buying w is $w - c^w_{min}$. Hence, naive consumers are willing to choose product v only if $v - \underline{f} \geq w - c^w_{min}$.

Second, we show that if $v - \underline{f} \geq w - c^w_{min}$, then the above is actually an equilibrium. To do so, it is sufficient to show that no firm prefers to unshroud product v. If a firm unshrouds product v, to attract consumers it must provide consumer value of at least as much as they would get from product w. Hence, a firm that unshrouds must provide value of at least $w - c^w_{min} \geq v - c_{min} \geq v - c_n$, which it cannot profitably do. $\square$

Proof of Proposition 5. Again, as in the proof of Proposition 3 we establish a slightly more general version of this proposition in which the maximum additional prices firms can impose differ across firms. Let $\overline{a}_n$ be the maximum additional price firm n can impose. We prove the statement of

Proposition 5 maintaining the assumption that the price floor is binding for all firms, i.e. $\underline{f} > c_n - \bar{a}_n$ for all n. Recall from the proof of Proposition 3 that with different maximal additional prices, a deceptive equilibrium exists for $\eta = 0$ if and only if for all firms n,

$$s_n(\underline{f} + \bar{a}_n - c_n) \geq v - c_n. \tag{6}$$

This proof has five steps.

(i): *No firm unshrouds the additional price with probability one.* If a firm unshrouds with probability one, all consumers become sophisticated and hence buy from the firm with the lowest total price $f + a$. Hence by the exact same argument as in Proposition 2, all consumers buy at a total price $f + a = c$ 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.

(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 $(\underline{f}, \bar{a}_n)$.

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

(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 *conditional* on firm n unshrouding; set $\hat{t}_n = 0$ in case firm n does not unshroud. Let $\hat{t} = \max_n \left\{\hat{t}_n\right\}$. 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). Also, by (iii), $\hat{t}$ is bounded from above and hence well-defined **!!! this sentence should be earlier. !!!**.

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 $v - c_n - \eta$, because the additional price is unshrouded by firm n and hence consumers never buy the product from firm n if $\hat{t} > v$. When firm n shrouds and offers $(\underline{f}, \overline{a}_n)$, however, its profits conditional on all other firms shrouding are at least $s_n(\underline{f} + \overline{a}_n - c_n)$. Thus, the equilibrium condition $s_n(\underline{f} + \overline{a}_n - c_n) \geq v - c_n$ implies that deviating by shrouding and offering $(\underline{f}, \overline{a}_n)$ is profitable—a contradiction.

Second, 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 η conditional on some other firm unshrouding—i.e. it loses money in expectation relative to shrouding and offering $(\underline{f}, \overline{a}_n)$. In addition, conditional on all other firms shrouding firm n earns less than 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.

(v): *All firms offer the contract* $(\underline{f}, \overline{a}_n)$ *with probability one.* By (iv), all firms choose to shroud with probability one. Hence, in equilibrium all firms charge an additional price $a = \overline{a}_n$ with probability one. By the exact same argument as in Proposition 2, all firms offer a base fee $f = \underline{f}$. □

Proof of Proposition 6. Part I. Consider a candidate equilibrium in which all firms other than n charge marginal cost in the superior-product market, firm n charges c^w_{min} in this market, all firms charge $\underline{f}, \overline{a}$ in the inferior-product market, and only firm n unshrouds. According to the tie-breaking assumption for the superior-product market, therefore, firm n attracts all superior-product consumers. Since lowering a attracts no new customers in the inferior-product market and raising f induces all inferior-product customers to buy from a rival, charging $\underline{f}, \overline{a}$ in the inferior product market is optimal for all firms. Furthermore, no firm $n' \neq n$ can profitably attract customers in the superior-product market, and thus both charging marginal cost and shrouding is a best response. It is also clearly optimal for firm n to charge c^w_{min} in the superior-product market, because in order to attract uneducated naive consumers firm n has to set a price of product w below its production cost. Furthermore, the decision of whether to shroud does not affect the behavior of sophisticated consumers, and hence we can focus on the profits earned from naive consumers. When shrouding, firm n earns

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

while when unshrouding it earns

$$(1-\kappa)[\lambda_n M+(1-\lambda_n)s_n(\underline{f}+\overline{a}-c_n)].$$

Rewriting shows that unshrouding is strictly optimal for firm n if $M > s_n(\underline{f}+\overline{a}-c_n)$ and hence under this condition the candidate equilibrium is indeed an equilibrium.

Part II. Consider the same candidate equilibrium as above except that firm n now shrouds. By analogous reasoning as in Part 1, this is indeed an equilibrium. □

Proof of Proposition 7. We define the term "exploitative equilibrium" as an equilibrium in which consumers buy the product from a firm setting $a_n > \hat{a}$ with positive probability.

Note that if $a_n \in (\hat{a}, \overline{a})$, increasing a_n to $\overline{a}$ does not change firm n's demand. Thus, without loss of generality we suppose no firm sets $a_n \in (\hat{a}, \overline{a})$ with positive probability in any equilibrium. By the same argument as in the proof of Proposition 3, whenever a firm sets $a_n = \overline{a}$, it charges the up-front price $\underline{f}$.

It is straightforward that an exploitative equilibrium in which all firms set $a_n = \overline{a}$ with probability one exists if Inequality (4) holds. Suppose that Inequality (4) does not hold for firm n' and an exploitative equilibrium exists. Then, in this equilibrium some firm sets $(\underline{f}, \overline{a})$ and consumers buy from the firm with positive probability. Note that in this case each firm can earn positive profits by setting $(\underline{f}, \overline{a})$. Let t_n be the supremum of firm n's total price distribution. Let $t \equiv \max_n t_n$. First, consider the case of $t > \underline{f}+\overline{a}$. Then, $t_n = t$ for all n; otherwise some firm earns zero profits when setting its total price above $t-\epsilon$ for some sufficiently small $\epsilon > 0$. Also, no firm can have an atom on the total price t. In this case, however, a firm's expected profit of setting the total price $(t-\epsilon, t)$ goes to zero as $\epsilon \to 0$ because t is bounded from above and the probability that the firm can get a positive market share by setting a price in that range goes to zero—a contradiction. Second, consider the case of $t = \underline{f}+\overline{a}$. Then, every firm sets the total price t with positive probability because consumers buy from the firm setting $(\underline{f}, \overline{a})$ with positive probability. Firm n', however, has an incentive to deviate from $t_{n'} = \underline{f} + \overline{a}$ and set $f_{n'} = \underline{f}$ and its additional price slightly below $\hat{a}$—a contradiction. Therefore, there is no exploitative equilibrium when Inequality (4) does not hold. $\square$