The Market for Deceptive Products*

Paul Heidhues
ESMT, CESifo, and CEPR

Botond Kőszegi
University of California, Berkeley

Takeshi Murooka
University of California, Berkeley

May 15, 2012

Abstract

We analyze conditions facilitating profitable deception and incentives for innovation in a competitive market for potentially deceptive products. Firms selling homogenous products simultaneously set a transparent "up-front price" and an "additional price," and decide whether to shroud the additional price from naive consumers. To model especially financial products such as banking and credit-card services, actively managed mutual funds, and non-traditional mortgages, we assume that there is a binding floor on the product's up-front price. 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. But perversely, if the product is socially wasteful, unshrouding would eliminate the industry, so in this case a profitable deceptive equilibrium always exists. 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. Regarding innovation decisions, because learning ways to charge higher additional prices increases the profits from shrouding and thereby lowers the motive to unshroud, a firm may have a strong incentive to make such exploitative innovations *and* have competitors copy them. In contrast, the incentive to make innovations that increase the product's value to consumers is zero or negative if other firms can copy the innovation, and even otherwise is strong only if the product is socially wasteful. **JEL Codes: D14, D18, D21**

*Kőszegi thanks the Center for Equitable Growth and the National Science Foundation (Award #0648659) for financial support. We are grateful to Özlem Bedre-Defolie, Stefano DellaVigna, Drew Fudenberg, Michael Grajek, Michael Grubb, Péter Kondor, Nicolas Melissas, and Antoinette Schoar for useful comments and feedback. We thank audiences at Carlos III, the 8th annual ACLE Competition & Regulation Meeting (Amsterdam), the Conference on Retail Financial Services after the Crisis (Frankfurt), EARIE (Stockholm), Federal Reserve Board of Governors, German Economist Abroad Meeting (Berlin), Industrieökonomische Ausschuss des VfS (Munich), LMU Munich, Pompeu Fabra, the Pros and Cons of Consumer Protection Conference (Stockholm), the University of Bayreuth, and the University of Copenhagen.

1 Introduction

In this paper, we investigate circumstances under which firms sell products by deceiving some consumers about the products' full cost. We build on the fundamental insight of Gabaix and Laibson (2006) that firms may shroud hidden fees to exploit naive consumers, but identify more specific product features and market forces that facilitate such deception, and focus in particular on deception that leads to positive equilibrium profits in seemingly competitive industries.[1] We show that if a floor prevents competing for profitable consumers by lowering the transparent up-front price, firms are often unwilling to instead compete by coming clean about hidden fees, as this would reveal to consumers that the product is expensive and hence lower industry demand. As a result, a profitable deceptive equilibrium can often be maintained by firms. Our theory identifies a perverse aspect of when this can happen: 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 products. 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. But if the product is socially wasteful, unshrouding would eliminate the industry, so in this case a profitable deceptive equilibrium always exists. In a market with multiple products, because a superior product both diverts sophisticated consumers and renders the inferior product socially wasteful in comparison, firms can always profitably sell the *inferior* product by deceiving consumers. We also identify a perverse incentive for socially wasteful innovation: because learning ways to charge consumers higher hidden fees increases the profits from shrouding and thereby lowers the motive to unshroud, a firm may

[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.

have a strong incentive to make such 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.

Section 2 introduces our model, in which firms are engaged in simultaneous-move price competition to sell a single homogenous product. Building on 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 them, naive consumers ignore the additional prices when making purchase decisions. In contrast to most existing work, however, we posit that there is a floor on the up-front price.[2] We investigate conditions under which a profitable shrouded-prices 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 which unshrouding has any cost, no matter how small, and all firms prefer it over an unshrouded-prices and hence zero-profit equilibrium.

The above model captures in a stylized way many markets, including primarily financial markets such as banking and credit-card services, actively managed mutual funds, and non-traditional mortgages. In each of these markets, there are costs—e.g. overdraft and other fees for bank accounts, fees, penalties, and long-term interest for credit cards, management fees for mutual funds, and penalties and future changes in monthly payments for mortgages—that many consumers may ignore when getting the product. And in each of these markets, the up-front price cannot drop (much) below \$0 without attracting "arbitrageurs" who just want to cash in on the negative price, creating a price floor.

Section 3 presents our basic results. As two benchmark cases, we show that if the price floor is not binding or consumers are sophisticated in that they observe and take into account additional prices, firms compete either on the up-front price or on the total price, so profitable deception cannot occur. If the price floor is binding and consumers are naive, however, profitable deception

[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. We briefly discuss this and other possible reasons for the price floor below. Grubb (2011), Ko (2011), and Armstrong and Vickers (2012) also analyze models with variants of our price-floor assumption.

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.

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 our model: 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. In our basic single-product model above, sophisticated consumers stay away from a socially wasteful industry, leaving our analysis unchanged. In a valuable industry, the motive to attract sophisticated consumers creates pressure to cut the additional price, thereby also creating pressure to unshroud the additional price. Hence, the higher the proportion of sophisticated consumers, the less likely it is that profitable deception occurs. In a multi-product market with a superior and an inferior product, however, often sophisticated and naive consumers selfseparate into buying the superior and inferior products, 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 analyze firms' incentives to invest in various forms of product innovation, contrasting innovations that increase the maximum additional price a firm can charge consumers ("exploitative" innovations) with those that increase the product's value to consumers ("value-increasing" innovations). In an equilibrium with profitable deception, increases in firms' additional prices cannot decrease the up-front price, so a firm always has an incentive to make an exploitative innovation and does not mind if competitors acquire the innovation. Furthermore, because an increase in the additional price competitors charge lowers their motive to unshroud by increasing their profits from shrouding, a firm may even *prefer* competitors to acquire an exploitative innovation. In contrast, a firm can never benefit from a value-increasing innovation that other firms acquire, and because such an innovation increases the motive to unshroud by raising the profits from expanding market share, a firm may even be willing to pay to avoid the innovation. These observations can help explain why firms in the financial market have been willing to make innovations that competitors can easily copy, and why these innovations often seem to have included exploitative features.

Finally, we explore value-increasing innovations that competitors cannot fully copy. Because this kind of innovation steals the consumers of any competitor who lags behind, in a socially valuable industry it leads to unshrouding and the loss of profits from deception. Hence, the incentive for such innovation in a socially valuable industry is weak, and is negative for small increases in value. In contrast, because unshrouding is not a concern in a socially wasteful industry—and an increase in market share increases profits—the incentive for the same kind of innovation in a socially wasteful industry is strong, and is non-trivial even for vanishingly small increases in value.

In Section 6, we turn to extensions and modifications of our framework. Most importantly, we consider a specification of consumer naivete in which consumers know all prices, but underestimate their own willingness to pay for an add-on. Although we assume firms cannot eliminate consumers' misprediction of their own behavior, we show that this alternative generates insights similar to those of our basic model.

In Section 7, we discuss the behavioral-economics and classical literature most closely related to our paper. While our model builds on a growing theoretical literature that 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, and has not analyzed exploitative innovation. Indeed, in previous models the competition for naive consumers returns all of the profits from hidden fees to consumers, so that these models cannot investigate market conditions that facilitate profitable deception. In addition, by the logic of previous theories a firm should gain little from innovations others can copy, so from the perspective of these theories much of the recent exploitative innovation—which was in easily copyable (and quickly copied) contract terms—is also a puzzle. Our framework is also reminiscent of classical switching-cost models in the basic implication that consumers can be induced to pay high additional prices once they buy a product. But with rational consumers such a model does not feature anything corresponding to the threat of unshrouding by competitors and (if consumers know or learn product values) does not predict the systematic sale of inferior products in competitive industries, and hence does not generate most of the results in this paper. We conclude in Section 8 with mentioning some policy implications of our findings, and by pointing out the importance of investigating more realistic forms of unshrouding than that postulated in our simple model.

2 Basic Model

2.1 Setup

In this section, we introduce our basic model of a market for potentially deceptive products. N 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. 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 total price $f_n + a_n$. We assume that the highest possible additional price firms can impose is $\overline{a} > 0$.[3] If consumers weakly prefer buying and are indifferent between all firms, firm n gets an exogenously given market share $s_n \in (0, 1)$. If consumers weakly prefer buying and are indifferent between a subset of firms, these firms split the market in proportion to s_n.[4]

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 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 facing an up-front price at the floor and an additional price of zero. In stating our results, we focus on identifying conditions for and properties of shrouded-prices 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 shrouded-prices equilibrium exists, however, it is more plausible than the unshrouded-prices equilibrium for a number of reasons. Most importantly, in Section 6.1 we show that in that case, the shrouded-prices equilibrium is the unique equilibrium in the variant

[3] We think of $\overline{a}$ as depending primarily on the nature of the industry and the willingness of consumers to pay for the product's services at various points ex post. In some industries, $\overline{a}$ can be very high: for instance, because a firm may be able to impose certain hidden fees or prices on a consumer multiple times, $\overline{a}$ could easily be greater than v.

[4] In later sections, we consider several alternatives to and extensions of the above basic framework. First and foremost, in Section 6 we 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. In addition, in Section 4 we consider also sophisticated consumers; in Section 6 we discuss the limited ways in which heterogeneity in v affects our conclusions; and in Section 5 we investigate whether and when firms want to invest to increase $\overline{a}$ or v. Finally, in Section 8 we informally discuss alternatives to our extreme assumption that a firm can instantaneously, fully, and costlessly educate all consumers, and also consider the possibility that consumer groups or other non-firm market participants might educate consumers.

of our model in which unshrouding carries a positive cost, no matter how small the cost is. In addition, a positive-profit shrouded-prices 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 shrouded-prices equilibrium.

Although our paper identifies conditions under which a deception-based positive-profit equilibrium exists, 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 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 an important message of our model.

2.2 Motivation for Key Assumptions

Our model has two key assumptions: that naive consumers might ignore the additional price when making purchase decisions, and that there is a floor on the up-front price. The former assumption is a simplified variant of many assumptions that have appeared in the literature on 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 a number of industries. In banking, credit-card, retail-investment, and mortgage services, for instance, consumers may be unaware of or underestimate many fees providers impose, or may underestimate the extent to which they will incur the fees.[5] For example, a consumer might not know about overdraft fees or might believe that she will never overdraft, and an investor might not realize how much of a premium she will pay in management fees when investing in an actively managed mutual fund rather than an index fund.[6] And while we interpret the additional prices primarily as financial prices, our model

[5] If consumers understand that they will have to pay some extra fees after purchase, but underestimate those fees, 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.

[6] Several empirical findings are consistent with the above general hypothesis. The Office of Fair Trading (2008) of the UK reports that 23% of consumers who incurred overdraft charges did not know about such charges beforehand, and some consumers seem to have underestimated the fees or thought they would not incur the fees. In 2006, about

applies 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.[7]

Our assumption of a price floor is supported both by theoretical arguments and by some empirical evidence on its implications for firm behavior. 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, based on a mean portfolio value of €105,356 Euros." 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 (2011) and Armstrong and Vickers (2012) impose a no-negative-prices constraint in the context of naive consumers, and Farrell and Klemperer (2007) discuss the same constraint

one-third of banks' revenue from bank accounts came from overdraft charges. More generally, Cruickshank (2000, pages 126-7) reports that most consumers do not know specific fees associated with their bank accounts, which can contribute to Stango and Zinman's (2009) finding that consumers incur many avoidable fees. Evidence by Wilcox (2003) and Barber, Odean and Zheng (2005) indicates that investors underweight operating expenses when choosing mutual funds. Agarwal, Driscoll, Gabaix and Laibson (2008) provide evidence that many credit-card consumers seem to not know or forget about various fees issuers impose. Cruickshank (2000, page 127-8) also documents that most consumers do not understand key mortgage features, and Woodward and Hall (2010) find that borrowers underestimate broker compensation. Hall (1997) reports that 97% of buyers do not know the price of a cartridge when buying their printer, and as revealed in a survey by UK's Office of Fair Trading, retailers believe 75% of consumers do not have an idea about printing costs. Finally, regulators are worried about the "bill shock" many mobile-phone consumers face when they run up charges they did not anticipate (Federal Communication Commission 2010).

There is also experimental evidence from other settings that individuals often ignore some components of a product's price. For instance, Chetty, Looney and Kroft (2009) show that consumers' shopping behavior depends on whether sales taxes are posted on a supermarket's shelf, and Hossain and Morgan (2006) find that bidders on eBay underweight shipping costs. These products are considerably simpler—and the additional prices are more obvious—than most of the ones we consider in this paper. The fact that consumers ignore additional prices even for these simple products suggests that the phenomenon could be quite widespread.

[7] 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 6.2 applies.

in the context of switching-cost models. As another possible microfoundation, Miao (2010) shows that if the additional price is that of an add-on, firms cannot distinguish old and new consumers, and buying the base product is a substitute for the add-on (such as for software, where updating the old version yields the new version), the price of the add-on can serve as a floor for the price of the base product. In addition, although this is difficult to capture formally, it might be the case that if prices dropped too low, consumers would become suspicious that "there is a catch" and not buy the product, effectively imposing a price floor.

Empirical evidence consistent with some implications of our model provides further support for our assumption of a price floor. First, as indicated by the significant resources they spend trying to compete for consumers, many firms selling only slightly differentiated products seem to value new business, but—in contrast to natural models without a price floor predicting that this will lead to lower prices—they compete using non-price methods. According to Evans and Schmalensee (2005), for instance, credit-card companies sent out 5 billion direct-mail solicitations in 2000, an average of approximately 3.9 solicitations per month for each household in the United States. Similarly, mutual funds often compete for consumers not by lowering prices, but by paying "independent financial advisors" to direct consumers to them. Second, Bar-Gill and Bubb (2012, forthcoming) find suggestive evidence that the 2009 Credit CARD Act had the intended effect of limiting over-the-limit and late-payment fees, while—in contrast to models without a price floor predicting that these losses are compensated by increases in other fees—it had no effect on annual fees, teaser rates, and other unregulated fees, and reduced banks' profits.

3 Profitable Deception

This section analyses 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:[8]

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

Since in a shrouded-prices 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 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, there is no incentive to unshroud, so that a shrouded-prices 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).

By Proposition 1, therefore, deception can occur in our model even if the price floor does not bind. Nevertheless, deception is likely to be less serious in this case than in the main cases we consider below: because firms have no incentive to invent new ways of charging additional prices that can be copied by competitors, they are more likely to limit themselves to obvious deception opportunities. Furthermore, 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 main cases below they do neither.

As a second benchmark, we consider the influence of a (binding or non-binding) price floor $\underline{f}$ when all consumers are sophisticated in that they observe additional prices and make purchase decisions based on the total price $f_n + a_n$. With such consumers, a floor in one component of the

[8] 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 affects neither firm profits nor consumer welfare.

price is insufficient to sustain positive profits:

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 full 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 additional price until profits are zero.

3.2 Naive Consumers with a Binding Price Floor

Taken together, Propositions 1 and 2 imply that for profitable shrouding 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.[9] To identify when a shrouded-prices 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 and undercut competitors. 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

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

and a shrouded-prices equilibrium exists if the same is true for all n. Furthermore, since $s_n < 1$, Condition (1) implies that $\underline{f} + \overline{a} > v$, so that in a shrouded-prices equilibrium consumers receive

[9] In Section 6.3, we argue that assuming this inequality for all n rather than some n is essentially without loss of generality.

negative utility. Proposition 3 states this result, and also says that if some firm violates Condition (1), 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* (1) *holds for all* n*, a shrouded-prices equilibrium exists. In any shrouded-prices equilibrium,* $f_n = \underline{f}$ *and* $a_n = \overline{a}$ *for all* n*, consumers receive negative utility, and firms earn positive profits. If Inequality* (1) *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. A firm that unshrouds and cuts the additional price by a little bit tells consumers not only that its product is cheaper than competitors', but also that the product is more expensive than they thought. If $\underline{f} + \overline{a} > v$—that is, if consumers' utility from buying the product at the current total price is negative—this surprise will lead consumers not to buy, so that the firm can attract consumers by unshrouding only if it cuts the additional price by a substantial margin. Since this may not be worth it, the firm may prefer not to unshroud.

Importantly, if purchasing the product at the total market price is optimal (i.e. if $\underline{f} + \overline{a} \leq v$), then—although they still face a negative surprise regarding the product's total cost of ownership—consumers are willing to buy from a firm that unshrouds and undercuts competitors' additional price by a little bit. Hence, in this case a shrouded-prices 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.[10]

Beyond showing that Condition (1) is sufficient for profitable deception to occur, Proposition 3 establishes the (technically more difficult) converse that it is also necessary: if the condition is violated for some firm, then there is no deception in equilibrium. This converse result is crucial for our arguments regarding industry concentration and innovation below. Intuitively, 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 sometimes all rivals set higher total prices when shrouding. For these high total prices, firms earn positive profits only when shrouding occurs, so that—very roughly speaking—similar arguments to those above imply that they set prices $\underline{f}, \overline{a}$. But then a firm that violates Condition (1) prefers to unshroud. Hence, we predict no deception whenever some firm violates Condition (1), and the classic Bertrand logic implies that firms earn zero profits.

3.3 Socially Valuable versus Socially Wasteful Products

We now use Proposition 3 to identify some circumstances under which deception does versus does not occur, distinguishing non-vanishingly socially valuable and socially wasteful products.[11]

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 (1) is positive and fixed. Hence, if a firm has a sufficiently low market share s_n when consumers are indifferent between firms, the firm violates Condition (1), and therefore a deceptive equilibrium does not exist. The intuition is simple: a firm that earns low profits from deception prefers to unshroud and capture the entire market, eliminating the possibility of profitable deception.

[10] Furthermore, if a shrouded-prices 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, in which the market share of firms other than the most efficient ones is zero.

[11] We do not discuss in-between cases in which the product is valuable to produce by some firms but not other firms. The implications below 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.

The above simple insight has two potentially important implications. First, it leads to interesting comparative statics with respect to industry concentration. For a relatively small number of firms N, it may be the case that Condition (1) holds for all firms and therefore a positive-profit shrouded-prices equilibrium exists. In this range, N has no effect on industry revenue at all—as N increases, the industry revenues are merely being divided among more firms. When there are sufficiently many firms, however, at least one will have s_n sufficiently low for Condition (1) to be violated, and therefore the industry undergoes a regime shift: firms switch to transparent pricing, revenues shrink and profits evaporate.[12]

Second, the insight that low profits from deception lead to unshrouding suggests a general competition-impairing feature in socially valuable industries. Given that firms make positive profits in a shrouded-prices equilibrium but zero profits in an unshrouded-prices equilibrium, to reduce the motive to deviate from the shrouded-prices equilibrium each firm wants to make sure competitors earn sufficient profits from shrouding. This feature leads to some socially inefficient innovation incentives we discuss in Section 5, and is likely to have other anticompetitive implications as well.

Socially wasteful product ($v < c_n$ for all n). In this case, the right-hand side of Condition (1) is negative while the left-hand side is positive. Hence, a shrouded-prices equilibrium exists regardless of the industry's concentration and other parameter values. This perverse result has a simple and compelling logic: the only reason consumers buy a socially wasteful product is that they are deceived about its total price, so that a firm cannot profitably sell the product by coming clean. Since a socially valuable industry with many firms must be transparent, our theory therefore says 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.

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 higher) increases in monthly payments.[13] While this mortgage may make sense for consumers who

[12] More precisely, if $N > (\underline{f}+\overline{a})/\epsilon$, 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 (1).

[13] 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://`

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 them to overborrow and get into financial trouble, the product might well have lowered many consumers' utility. 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 bad for consumers and socially wasteful.

4 Sophisticated Consumers

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: because sophisticated consumers cannot be fooled by shrouded prices into mistakenly buying a product, they increase firms' incentive to unshroud and compete in the additional price. We go on to show in Section 4.2, however, that if there is a superior and transparent alternative product, this product often not only relieves sophisticated consumers' pressure to unshroud, but in fact *guarantees* that firms can exploit naive consumers.

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

4.1 Sophisticated Consumers in Our Basic Model

If sophisticated consumers buy from any firm in equilibrium, then unshrouding must occur because a firm with the lowest total price prefers to unshroud and attract all consumers. Hence, in any

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).

shrouded-prices equilibrium sophisticated consumers do not buy the product. To derive the condition for such an equilibrium to exist, 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. As the analogue of Inequality (1), this is unprofitable if

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

Hence, a shrouded-prices equilibrium exists if and only if Inequality (2) holds for all n.

Condition (2) for the existence of a shrouded-prices 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 shrouded-prices 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 shrouded-prices 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 we have discussed above, 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, suppose that in addition to the product we have assumed throughout the paper, each firm also has a transparent product with value $w > 0$, where firm n's cost of producing product w is c_n^w. We assume that $\min_n\{c_n^w\} = c_{min}^w > 0$, and that there are at least two firms whose cost of producing product w is c^w_{min}. Crucially, we posit that product w is socially valuable ($w - c^w_{min} > 0$), and it is not inferior to product v: $w - c^w_{min} \geq v - c_{min}$. 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. Then:

Proposition 4 (Separation of Naive and Sophisticated Consumers)**.** *There exists an equilibrium in which each firm shrouds the additional price of product v, naive consumers buy the shrouded product v, and sophisticated consumers buy the transparent product w, if and only if $v - \underline{f} \geq w - c^w_{min}$. 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 $v - \underline{f} \geq w - c^w_{min}$, a positive-profit equilibrium in which naive consumers are deceived always exists. In this equilibrium, 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 shrouded product is costly but naive consumers believe it is a better value, in equilibrium the two types of consumers separate. Second, if a firm unshrouded the additional price of the inferior product, consumers would immediately realize that the other product is a better deal, 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 naive consumers may realize that a debit card serves essentially the same functions as a credit card, they may prefer a credit card because they falsely believe that its perks (e.g. cash-back bonuses) make it a better deal.

By Proposition 4, therefore, 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. To ensure the existence of such a product and hence maintain a positive-profit deceptive equilibrium, firms often have strong incentives to invent a superior *or* inferior alternative product to the one already in the market. For instance, firms might want to create a superior product even if it means losing money on that product.

As an example, because few mutual-fund managers can persistently outperform the market (Carhart 1997, Kosowski, Timmermann, Wermers and White 2006), the service many or most managers provide—actively investing instead of tracking an index—is not worth the high fees they charge. As a result, index funds are superior to most managed funds, and for this reason 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 product w is superior—that is, $w - c_{min}^{w} \geq v - \overline{a} - c_{min}$—Proposition 4 and the logic behind it survive unchanged.

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 always exists. 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.[14]

Going beyond the setting of our model, the insights above have an immediate implication for the marketing of superior and inferior products that contrasts with classical views of advertising. 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. This pushing can take a number of different forms. First, firms may inform consumers unaware of the inferior product of the product's existence, yet not do the same for the superior product. Second, firms may pay intermediaries to convince consumers to buy the inferior product. Third, firms may make costly (real or perceived) improvements to the inferior product to make it more attractive to consumers.[15] Indeed, Anagol, Cole and Sarkar (2011) and Mullainathan, Nöth and Schoar (2011) document that intermediaries tend to disproportionately push inferior products in the mutual-fund and life-insurance markets, respectively, and do so because they receive higher commissions from firms. And while the possibility that firms push only the socially inferior product arises naturally in our model, it is inconsistent with equilibrium in a model with rational consumers: a rational consumer would know that the firm is pushing the inferior product, and realizing that an inferior product a firm would like to sell must generate lower consumer surplus than the alternative, she would prefer to buy the alternative.

[14] The above logic does not require that $w > v$ for all consumers, only that this is the case for a positive fraction of naive consumers. If $v > w$ for all naive consumers, then there can also be an equilibrium in which both additional prices are shrouded, naive consumers buy product v, and sophisticated consumers buy product w. In this equilibrium, firms do not unshroud the additional price of product w, but sophisticated consumers nevertheless understand that it is the superior product.

[15] To see one formalization of the above ideas, consider a variant of our two-product model in which some naive consumers are unaware of the existence of the inferior and some of the superior product, while sophisticated consumers are aware of both products. Each type of consumer acts in the way we have defined above, given the product she is aware of. Firms play the following game. First, firms simultaneously choose the prices and whether to unshroud (as above), and each consumer chooses a firm to buy from. When a consumer approaches a firm to buy a particular product, the firm can point out the existence of the other product at a cost $\epsilon > 0$. If the firm does so, the consumer considers both products at all firms, but if she does not find what she thinks is a better deal, she buys from this firm. Then, if ϵ is sufficiently small, there is a profitable deceptive equilibrium in which firms shroud the additional price of the inferior product, price the superior product at marginal cost, inform consumers of the inferior product, but do not inform consumers of the superior product.

5 Research and Development Incentives

While so far we have taken the product firms sell as exogenous, in most markets motivating our analysis firms can engage in innovation to change their product. Within our reduced-form model, this product innovation can have two effects: changing the extent to which firms can unexpectedly charge naive consumers ($\overline{a}$) or changing the product's value to consumers (v). In this section, we identify a number of socially inefficient incentives with regard to such exploitative and value-increasing innovation. We show that—independently of whether the product is socially valuable and whether competitors can copy the innovation—a firm's incentive to make an exploitative innovation is often strong, but its incentive to make a value-increasing innovation is strong only if the product is socially wasteful and others cannot copy the innovation. Combining these insights with our previous results, our theory predicts some extremely inefficient behaviors: it says, for instance, that firms will market a socially wasteful product, spend lots of money devising new hidden fees for this product, and waste yet more money trying to improve the product by tiny amounts.

To identify innovation incentives in a transparent manner, we assume that only one firm, firm 1, can make innovation investments, and analyze its investment decisions separately for exploitative and value-increasing innovation. In each case, we modify the pricing game above by assuming that there is an initial stage in which firm 1 chooses whether or not to invest, with all firms observing its decision. An exploitative innovation costs I_a and increases the maximum additional price firm n can charge by Δa_n, where $0 \leq \Delta a_n \leq \Delta a_1$. This formulation allows us to consider the two extreme cases often studied in the literature, appropriable innovations (which competitors cannot copy: $\Delta a_n = 0$ for all $n > 1$) and non-appropriable innovations (which competitors can fully copy: $\Delta a_n = \Delta a_1 > 0$ for all $n > 1$), as well as in-between cases. Analogously, a value-increasing innovation costs I_v and increases consumers' valuation of firm n's product by Δv_n, where $0 \leq \Delta v_n \leq \Delta v_1$.[16] Following our equilibrium-selection arguments in Section 2, we assume that whenever an equilibrium in which prices are shrouded with probability 1 exists in the pricing subgame, firms play such a continuation

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

equilibrium after firm 1's innovation decision.[17] For simplicity, we also impose throughout that the price floor is originally binding for firm 1 ($\underline{f} + \overline{a} > c_1$). We look for the maximum investment costs I_a^*, I_v^* below which firm 1 is willing to make the investment of each type.

Our key results in this section derive from considering how firm 1's investment affects other firms' willingness to go along with deceiving consumers. To refer to this willingness, we introduce:

Shrouding Condition. Firm n satisfies the Shrouding Condition at $(\hat{a}_n, \hat{v}_n)$ if $s_n(\underline{f} + \hat{a}_n - c_n) \geq \hat{v}_n - c_n$.

Proposition 5 states our results for non-appropriable innovations:

Proposition 5 (Non-Appropriable Innovations)**.**

I. (Exploitative.) Suppose $\Delta a_n = \Delta a > 0$ for all n. If all firms n satisfy the Shrouding Condition at $(\overline{a} + \Delta a_n, v)$, then $I_a^* \geq s_1 \Delta a$. If in addition some firm does not satisfy the Shrouding Condition at $(\overline{a}, v)$, then $I_a^* = s_1(\underline{f} + \overline{a} + \Delta a - c_1) > s_1 \Delta a$.

II. (Value-Increasing.) Suppose $\Delta v_n = \Delta v > 0$ for all n. Then, $I_v^* \leq 0$.

Part I of Proposition 5 says that if firms can maintain a deceptive equilibrium following the innovation, firm 1 is willing to spend resources on—socially clearly wasteful—non-appropriable exploitative innovation. In this situation, increasing the additional price from $\overline{a}$ to $\overline{a} + \Delta a$ cannot lead to a decrease in the up-front price, so the innovation increases firm 1's profit by at least $s_1 \Delta a$. Going further, if in addition firms cannot maintain a deceptive equilibrium without the innovation, the innovation increases firm 1's profits even more by enabling profitable deception in the industry. In this case, firm 1's willingness to pay for innovation is equal to its full post-innovation profits—a potentially huge incentive to innovate.

Part II of Proposition 5 says that in contrast to socially-wasteful exploitative innovation, firm 1 is not willing to spend on—socially often valuable—non-appropriable value-increasing innovation. Because such innovation can increase neither one's market share nor one's markup, firm 1 has

[17] In Proposition 7, only firm 1 earns positive profits in a shrouded-prices equilibrium following innovation, so our less important equilibrium-selection argument in Section 2—that all firms prefer the shrouded-prices equilibrium to an unshrouded-prices equilibrium—does not apply with the same force as previously. But our more important argument applies with equal force: because Proposition 7 deals with a socially wasteful industry, unshrouding is suboptimal whenever it carries a positive cost, so that the shrouded-prices equilibrium is the unique continuation equilibrium with positive unshrouding costs.

no incentive to invest in it. More interestingly, it may be the case that firms can maintain a deceptive equilibrium without but not with the innovation, so that firm 1's willingness to pay for the innovation is negative. Intuitively, because an increase in v does not affect profits when firms shroud but increases the profits a firm can gain from expanding market share, it increases the motive to unshroud. As a result, firm 1 may be willing to spend money to avoid an increase in v.

To sharpen the intuitions from Proposition 5 on a firm's incentive to invest in exploitative innovation, we compare these incentives for non-appropriable and appropriable innovations:

Proposition 6 (Appropriability of Exploitative Innovation)**.** *I_a^* is weakly greater if the innovation is non-appropriable ($\Delta a_n = \Delta a > 0$ for all n) than if it is only partially copyable ($\Delta a_1 = \Delta a \geq \Delta a_n \geq 0$ for all $n \neq 1$), and it is strictly greater if the Shrouding Condition holds at $(\overline{a} + \Delta a, v)$ for all n but fails at $(\overline{a} + \Delta a_n, v)$ for some n.*

Proposition 6 says that firm 1 has a weakly *greater* incentive to engage in exploitative innovation when other firms can copy its innovation than when they cannot—and, equivalently, firm 1 weakly prefers others to obtain its innovation. As above, an increase in a competitor's additional price does not lead to greater competition in the up-front price, so it never lowers firm 1's profit. Moreover, a competitor who is not very good at imposing additional prices gains little from deception and hence may want to deviate from it, threatening the deceptive equilibrium and thereby firm 1's profits. To eliminate such a threat, firm 1 would like to teach this competitor how to better exploit consumers.[18]

The message of Proposition 5 that firms might be willing to make investments in non-appropriable innovations, and that such innovations are likely to be exploitative rather than socially valuable, seems consistent with how consumer financial products have developed recently. As we have argued above, many consumers likely underappreciate a number of future fees and other payments associated with credit cards, bank accounts, and non-traditional mortgages. And as has been argued by Heidhues and Kőszegi (2010), Grubb (2011), and other researchers, firms exacerbate these

[18] An important caveat to Proposition 6 is that—while willing to share with competitors already in the marketplace—a firm often prefers not to share an exploitative innovation with potential entrants. The ability to charge higher additional prices can make participation in the market more attractive and hence induce additional entry, reducing the innovator's market share.

mispredictions with carefully designed contract features—such as teaser rates, high fees for certain patterns of product use, and difficult-to-understand payment schedules involving large future payments—whose main purpose is likely to hide products' total use cost to consumers. Furthermore, not only are such exploitative contract innovations extremely easy to copy (and in fact quickly copied), in some instances firms seem—consistent with Proposition 6—positively willing to share them with each other.[19]

To complete our analysis of innovation decisions, we consider fully or partially appropriable value-increasing innovations. For this part of our analysis, we assume that no firm charges a total price below its marginal cost, and that consumers go to a highest-quality firm when indifferent.[20] We begin our analysis with socially wasteful industries, where, to make our result easier to state and understand, we consider only innovations that make firm 1's product better than that of any competitor.

Proposition 7 (Value-Increasing Innovation in Socially Wasteful Industries). *Suppose $v+\Delta v_n < c_n$ for all n, and $\Delta v_1 > \Delta v_n$ for all $n \neq 1$. Then,*

$$I_v^* = \left[(1-s_1)(\underline{f}+\overline{a}-c_1)\right] + \left[\Delta v_1 - \max_{n\neq 1} \Delta v_n\right] > 0. \tag{3}$$

Proposition 7 implies that firm 1's willingness to pay for fully or partially appropriable value-increasing innovations in a socially wasteful industry is quite high: it is greater than in the corresponding classical setting (where it would clearly be zero), it is greater than the increase in the relative value of firm 1's product ($\Delta v_1 - \max_{n\neq 1} \Delta v_n$), and (because I_v^* is bounded away from zero) it is non-trivial even for vanishingly small product improvements. Firm 1's willingness to

[19] For example, Argus is an information-exchange service that collects individual-level account data from credit-card issuers and, based on this data, relays information on current practices to other issuers. The information Argus collects includes fee assessment practices, strategies for balance generation, financial performance, and payment behavior. Argus emphasizes that it has detailed information on "virtually every US consumer credit card." (See `http://www.argusinformation.com/eng/our-services/syndicated-studies/credit-card-payment-study/us-credit-card-payments-study/`.) For any issuer who participates, any innovation is essentially a non-appropriable innovation. Proposition 5 explains why a participating firm makes innovations, and although there may be other reasons, Proposition 6 provides a strategic reason for why a firm that is interested in developing new exploitative practices in the credit market is willing to join Argus in the first place.

[20] The first assumption is typical in classic Bertrand-competition models. The second assumption allows a firm to price a lower-quality competitor out of the market by offering the same deal—something it could do anyway by offering a minimally better deal—ensuring the existence of equilibrium.

pay, I_v^*, derives from two sources. First, as captured in the first term in Equation (3), the innovation attracts the consumers of all competitors to firm 1, and firm 1 benefits from this even at pre-innovation market prices. Second, as captured in the second term in Equation (3), because the innovation makes firm 1's product better than the best alternative, firm 1 can increase the up-front price without losing consumers, further increasing its profits. Although firm 1 makes these extra profits by pricing competitors out of the market, with the industry being socially wasteful competitors have no incentive to respond by unshrouding.

We next consider socially valuable industries. In this setting, a useful benchmark for a firm's innovation incentives is based on what could be called its "relative advantage"—its ability to deliver value over and above its competitors:

$$RA_n(v_1, \ldots, v_N) \equiv \max\{(v_n - c_n) - \max_{k \neq n}(v_k - c_k), 0\}.$$

As is recognized in the literature, in classic Bertrand competition a firm earns its relative advantage in any equilibrium in which firms do not price below marginal cost.[21] Because absent innovation firm 1's relative advantage is zero, this implies that in a classical setting its willingness to pay for an innovation would be $RA_1(v + \Delta v_1, \ldots, v + \Delta v_N)$, the relative advantage it establishes through the innovation.

Proposition 8 (Value-Increasing Innovation in Socially Valuable Industries)**.** *Suppose $v > c_n$ for all n, $\Delta v_1 > \Delta v_n$ for some $n \neq 1$, and the Shrouding Condition holds for all firms n at $(v, \overline{a})$. Then, $I_v^* = RA_1(v + \Delta v_1, \ldots v + \Delta v_N) - s_1(\underline{f} + \overline{a} - c_1) < \big[(1 - s_1)(\underline{f} + \overline{a} - c_1)\big] + [\Delta v_1 - \max_{n \neq 1} \Delta v_n]$. If $\Delta v_1 < s_1(\underline{f} + \overline{a} - c_1)$, then $I_v^* < 0$.*

Proposition 8 implies that firm 1's willingness to pay for partially appropriable value-increasing innovation in a socially valuable deceptive industry is quite small: it is lower than in the corresponding classical setting, it is lower than in a socially wasteful industry with the same relative product values, and it is negative for non-substantial improvements. Any innovation that improves

[21] For example, if firms' products are homogeneous, the relative advantage of all but the lowest-cost firms is zero, and the relative advantage of a lowest-cost firm is the difference between the second-lowest cost and its cost. In addition, a lowest-cost firm earns this relative advantage in an equilibrium in which it charges the second-lowest cost, with lowest-cost firms getting the entire market.

firm 1's product above that of firm n must lead to unshrouding, as otherwise firm n would have zero market share and hence would choose to unshroud itself. Thus, innovation leads firm 1 to lose its positive profits from deception. This loss dampens firm 1's incentive to innovate, and for small improvements—which generate only a small relative advantage—the incentive is negative.

6 Extensions and Modifications

To demonstrate some robustness of our findings, in this section we discuss various extensions and modifications of our framework.

6.1 A Model with Costly Unshrouding

In this section, we provide one possible justification for our presumption that firms play a shrouded-prices equilibrium whenever it exists: that if it exists, 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. We consider the same game as in Section 2, except that we now assume unshrouding costs a firm $\eta \geq 0$.

Proposition 9. *Fix all parameters other than η, and suppose a shrouded-prices equilibrium exists for $\eta = 0$. Then, if $\eta > 0$, there exists a unique equilibrium, and in this equilibrium all firms shroud and offer $(\underline{f}, \overline{a})$ with probability one.*

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 would have an incentive to shroud in this situation with zero unshrouding cost—which is exactly the condition for a shrouded-prices equilibrium to exist—then it strictly prefers to shroud with a positive unshrouding cost.

6.2 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. Similarly, a mobile-phone consumer may know that going beyond her included minutes, or making calls or texting during a trip, can be expensive, but underestimate her use of these add-on services.[22]

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 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 $v - \hat{a}$.[23] 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—while highlighting an otherwise hidden price component may be relatively easy—convincingly explaining to a consumer how she herself will behave is very difficult. Indeed, a consumer may be presented with and readily believe information about how the average consumer behaves, but still believe that this does not apply to her.

[22] Consistent with the first 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. And consistent with the second example, Grubb (2009) documents that many mobile-phone consumers choosing contracts with high fees for high usage would have been better off choosing a plan with lower fees for high usage—while few consumers make the opposite mistake.

[23] 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.

Proposition 10 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 prices are not shrouded in the current model, we refer to such an equilibrium as an "exploitative equilibrium" rather than a shrouded-prices equilibrium.

Proposition 10 (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*

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

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. But 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.[24]

Given the similarity of Conditions (1) and (4), the implications of Proposition 10—as well as those of introducing sophisticated consumers or allowing a firm to innovate—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 it 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

[24] 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.

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. Indeed, the number of firms in the credit-card market has increased drastically over time, and although there are sharp disagreements on the extent, interpretation, and economic implications of this finding, firms continue to make seemingly large profits from interest charges and fees on existing consumers (Ausubel 1991, Bar-Gill 2004, Evans and Schmalensee 2005).[25]

6.3 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 shrouded-prices 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 shrouded-prices 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 shrouded-prices 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

[25] As emphasized above, this does not necessarily mean that credit-card issuers earn positive profits net of entry, marketing, product development, and other costs. Nevertheless, a model of price competition in barely-differentiated products seems to apply well to the credit-card market *after* entry and marketing costs have been paid and consumers have been reached by multiple firms with very high probability, but consumers have not signed on yet. Hence, from the perspective of prior research it is puzzling how credit-card issuers can sustain positive profits at this stage.

least one firm would choose to unshroud, eliminating the shrouded-prices equilibrium.

Now consider 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 purchase 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.

All of our main results survive unchanged if we define sophisticated consumers not as those who observe additional prices, but as those who understand firms' equilibrium pricing behavior.[26] Intuitively, such "strategically sophisticated" consumers understand that since they cannot observe shrouded additional prices, firms have an incentive to set such prices to be the maximum $\overline{a}$. As a result, if there are only strategically sophisticated consumers, unshrouding and undercutting competitors does not lower industry demand, so that a positive-profit shrouded-prices equilibrium cannot be maintained (Proposition 2). Similarly, as do sophisticated consumers in our basic model, strategically sophisticated consumers do not buy an overpriced product, leading to the same condition for when a shrouded-prices equilibrium exists in a market with both sophisticated and naive consumers (Condition (2)), and for when separation of sophisticated and naive consumers occurs (Proposition 4).

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 $\overline{a}_n$.[27] This assumption substantively modifies only Proposition 1: 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 shrouded-prices

[26] Technically, to model such strategically sophisticated consumers, we cannot use the approach above of defining Nash equilibrium only for firms. Instead, it is necessary to think of strategically sophisticated consumers as players, and look for a natural extension of perfect Bayesian equilibrium. In this equilibrium, firms maximize profits given others' behavior, a naive consumer acts as we have defined above, and a strategically sophisticated consumer maximizes her own welfare given her (in equilibrium correct) predictions of firm behavior.

[27] In fact, our proof of Proposition 3 in the appendix allows for this possibility.

equilibrium, but the firms with the lowest $c_n - \overline{a}_n$. Thus, in a shrouded-prices equilibrium of a socially wasteful industry now both productive and allocative efficiency may fail to hold.

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 (1) 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 Condition (1) 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.

7 Related Theoretical Literature

In this section, we discuss theories closely related to our paper. Although we identify other differences below, the most important one is that the previous literature does not address in any detail the central issues we consider: the role of socially wasteful and inferior products in maintaining profitable deception, and the incentives for exploitative innovation in a market for profitable deceptive products.

The model by Gabaix and Laibson (2006) is both the most closely related to ours and a foundation for it. In their model consumers buy a base good and then have the option to purchase an add-on whose price might be shrouded by firms. Sophisticated consumers anticipate the equilibrium add-on price, and they take costly steps to avoid the add-on if they believe its price will be high. Naive consumers do not anticipate high add-on prices and therefore do not take steps to avoid the add-on, ending up having to purchase it. 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 do not buy the add-on).[28] Although the precise

[28] 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

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.[29]

Grubb (2011) analyzes what in our language could be called the regulated unshrouding of additional prices, and develops a set of results complementary to ours. He 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.

In Spiegler (2006a), products have multiple price (or value) components, and firms compete for consumers who choose a product based on a single random price component. In equilibrium, firms randomize their prices, trying to attract consumers with a low-price component and cash in on high-price components. As the number of firms increases, prices in the high range of a firm's distribution become less effective at attracting consumers, so that firms *increase* these prices to at least make more money on consumers they do attract. As a result, profits remain bounded away from zero. Because in Spiegler's model consumers do not pay full attention to the price components, in our terminology these components correspond best to additional prices. Hence, Spiegler's model can be thought of as a microfoundation for how firms can sustain positive profits from additional prices even when consumers do not fully ignore those prices.

Spiegler (2006b) analyzes a model of the market for "quacks" (producers who provide no value

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.

[29] We have reconsidered our main questions in a model based on Gabaix and Laibson's framework combining naive and sophisticated consumers who can avoid the add-on, imposing that products are perfect substitutes. Our results that a firm does not want to unshroud the additional price of a socially wasteful product, and that a superior transparent product can serve to separate sophisticated and naive consumers, also hold in their setting. In contrast, because in their original model the unshrouding decision is driven solely by the comparison of the gains from trading the add-on and the cross-subsidy sophisticated consumers receive, the number of firms has no effect on whether shrouding occurs in equilibrium. In addition, because profits in their model are zero, firms never make non-appropriable investments.

relative to the outside option) in which consumers overinfer from a signal received about a quack and hence are too prone to choose a quack for whom they have observed a good signal. Consumers' mistaken inference creates a form of product differentiation that allows firms to raise prices and profits. Spiegler's model applies well, for instance, to the market for managed mutual funds, where consumers may overinfer from funds' recent performance.

Our theory, and those discussed above, also build on 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), Laibson and Yariv (2007), Grubb (2009), and Heidhues and Kőszegi (2010).

At a broad level, when consumers must pay classical search costs to find out prices (or product features), one can think of the prices as being partly shrouded, and firms' attempts to increase search costs can be interpreted as a type of exploitative innovation. Indeed, Stahl (1989) shows that an increase in search costs increases firm profits, and Ellison and Wolitzky (2009) establish that as a result, it can be individually rational for firms to increase search costs. While we believe that search costs are extremely important in the markets we consider, by themselves they do not seem to fully explain the behavior of firms in these markets. In a positive-profit equilibrium of a rational search-cost model along the lines of Stahl (1989) and Ellison and Wolitzky (2009), there cannot be a cheap binding price announcement that firms could make to at least some set of consumers who might otherwise end up searching and buying from other firms, and that firms are not making with high probability.[30] Yet in many markets we consider, there seem to be very cheap such announcements that firms are not making.[31] In addition, although we have no precise

[30] To see the intuition, suppose that announcements are costless. Consider a candidate positive-profit equilibrium with no announcement, and take the lowest price any firm sets in this equilibrium. At this price, the firm setting it makes its equilibrium—and hence positive—profits. By setting this price and making the announcement, the firm would increase market share and therefore profits, contradicting equilibrium. See Grossman and Hart (1980) for details. Baye, Morgan and Scholten (2006) illustrate that if announcements are costly but sufficiently cheap, there exists a mixed-strategy equilibrium in which firms randomize the announcement decisions. As announcement costs approach zero, the probability of no announcement also approaches zero.

[31] As discussed below, for instance, credit-card issuers could promise never to charge late fees. For a completely different example, consider the practice of internet computer-parts retailers discussed by Ellison and Ellison (2009), whereby consumers were charged exorbitant shipping fees after navigating to a retailer's homepage from a price comparison site. A firm could easily have announced its shipping charges or total price on the price-comparison cite.

empirical evidence, it does not seem that firms are playing the mixed-strategy pricing equilibrium predicted by these models.

Our theory shares one basic premise with the large literature on switching costs: that consumers can be induced to pay high additional fees once they buy a product. 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.[32] First, such a model does not predict a consideration analogous to the threat of unshrouding by competitors, and hence it does not generate most of our results regarding the effect of entry and the incentives for innovation. Second, if consumers know or learn product attributes, a rational switching-cost model seems to predict 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.[33]

8 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. 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.[34] If Condition

[32] The assumption of no commitment seems implausible in many industries motivating our analysis: a credit-card issuer, for instance, can easily include in its contract that it will never charge late fees. If so, in a rational switching-cost model an equilibrium in which firms charge high additional prices and make positive profits does not exist: a firm making positive profits on a consumer always has an incentive to commit to slightly lower prices than the competition.

[33] 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.

[34] 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

(1) 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. Intuitively, since consumers remain profitable, firms do not raise f in response to the policy. In addition, a decrease in the additional price can lead to some firm violating Condition (1), in which case the market becomes transparent, prices drop further, and productive efficiency obtains. Our analysis in Section 5, however, implies that in the latter case the policy greatly increases firms' incentive to make new exploitative innovations, and hence may have little net effect.

Relatedly, the regime shift from deceptive to transparent pricing predicted by our model for socially valuable products identifies a novel interaction between competition and consumer-protection policies. Roughly speaking, classical merger analysis in the US and Europe attempts to predict how changes in market structure affect consumer surplus and welfare. Our model highlights a potential change in industry conduct as the number of firms decreases beyond a critical threshold: the focus on inventing hidden fees and unexpected charges to the detriment of customers. This threshold is reached slower if, for instance due to consumer-protection policies mentioned above, the maximum additional price $\overline{a}$ is lower.

Throughout this paper, we have assumed that if a firm unshrouds, all firms' pricing schemes become known to all consumers. An important agenda for future research is to investigate the effect of more realistic unshrouding technologies. A simple first step is to assume that a firm can only educate a fraction of consumers, or that educating a consumer has some costs. In this case, the conditions for profitable deception to occur will of course be weaker, but qualitatively similar insights to those we find will still hold. Another alternative assumption is that firms can unshroud some fees but not all fees; an interesting question is then whether firms are more likely to unshroud small or large fees. Similarly, a firm may be able to unshroud its own fees but not other firms'

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 $\overline{a}$.

fees.[35]

The implications of more realistic unshrouding technologies are important not only for understanding firm behavior, but also for analyzing the potential impact of education campaigns by a social planner or consumer group. Literally interpreting our assumption that unshrouding products' additional prices is costless, a single consumer group or well-intentioned individual could eliminate profitable deception. If education is not so simple, however, this may no longer be the case. For example, if educating consumers is costly, it is unclear how a consumer group would finance an education 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 much 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.

In addition, given our emphasis that exploitation requires innovation, it would be interesting to investigate the dynamics of how exploitation appears and spreads in an industry. One possible scenario suggested by our theory is the following. The industry is initially in a non-deceptive situation (e.g. offering only 30-year fixed-rate mortgages). Then, one firm invents and starts offering a product with shroudable features (e.g. an Option ARM), and because neither consumers nor competitors were aware of this product, it starts off being shrouded. At this point, competitors must decide whether to unshroud the product or to adopt it in their product line. Our theory suggests that competitors' preference is to adopt the deceptive product.

Finally, this paper does not address in detail the policy questions of how a planner can detect whether consumers misperceive product features, and what the planner can do to make markets less deceptive. An important agenda for future research is to find ways to identify consumer mistakes from market data available to regulators and other observers for many markets, so that checking whether a particular market is deceptive does not require special data or expensive methods.

[35] In some situations, this ability can be sufficient to eliminate the shrouded-prices equilibrium. Suppose that the monthly base price of a mobile-phone contract is \$50 and additional charges amount to \$50, but consumers believe they will only have to pay \$10 in additional fees. Then, a provider that can credibly offer a package with a \$59 monthly fee and no additional charges attracts all consumers, even without revealing to consumers the charges other providers impose.

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.

Bar-Gill, Oren, "Seduction by Plastic," *Northwestern University Law Review*, 2004, *98* (4), 1373–1434.

____ **and Ryan Bubb**, "Credit Card Pricing: The CARD Act and Beyond," *Cornell Law Review*, 2012, forthcoming, *97.*

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.

Baye, Michael R., John Morgan, and Patrick Scholten, "Information, Search, and Price Dispersion," in Terrence Hendershott, ed., *Handbook on Economics and Information Systems*, Vol. 1, Elsevier, 2006, chapter 6, pp. 323–376.

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 and Alexander Wolitzky, "A Search Cost Model of Obfuscation," August 2009. NBER Working Paper #15237.

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

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

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.

Grossman, Sanford J. and Oliver D. Hart, "Disclosure Laws and Takeover Bids," *Journal of Finance*, 1980, *35* (2), 323–334.

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

——, "Bill Shock: Inattention and Price-Posting Regulation," 2011. 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`.

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.

Miao, Chun-Hui, "Consumer Myopia, Standardization and Aftermarket Monopolization," *European Economic Review*, 2010, *54* (7), 931–946.

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.

Stahl, Dale O., "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 1989, *79* (4), 700–712.

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," 2010. Working Paper, Stanford University.

A Proofs

Proof of Proposition 1. First, we establish properties of any shrouded-prices 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. $\square$

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

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. This is useful for Proposition 6 on firms' incentives to change this additional price. Let $\overline{a}_n$ be the maximum additional price firm n can impose. We prove the statement of Proposition 3 with Inequality (1) 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 shrouded-prices 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 (5) holds for all n, then there is a shrouded-prices 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 shrouded-prices 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 (5) 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 (5). 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 (5), because conditional on others shrouding it would earn $v - c_n$. □

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 shrouded-prices 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. □

Proof of Proposition 5. We first prove Case I. In the subgame following an innovation by firm 1, the Shrouding Condition holds for all firms by assumption, and thus firm 1 earns $s_1(\underline{f} + \overline{a} + \Delta a - c_1)$ in this case. In the subgame in which firm 1 did not innovate, firm 1 earns $s_1(\underline{f} + \overline{a} - c_1)$ if the Shrouding condition holds for all firms and zero otherwise. In the former case the innovation increases firm 1's profits by $s_1\Delta a$, in the latter case by $s_1(\underline{f}+\overline{a}+\Delta a-c_1)$, which is strictly greater than $s_1\Delta a$ because $\underline{f}+\overline{a}>c_1$.

We now prove Case II. Firm 1 earns zero profits in the pricing subgame whenever some firm violates the Shrouding Condition. If all firms satisfy the Shrouding Condition, firm 1 earns $s_1(\underline{f}+\overline{a}-c_1)$ which is positive and independent of v. The result, hence, follows from the fact that an increase in v either does not affect whether the Shrouding Condition holds or leads to a violation of the Shrouding Condition for some firm. $\square$

Proof of Proposition 6. Firm 1 earns $s_1(\underline{f}+\overline{a}+\Delta a-c_1)$ in the subgame following its innovation if the Shrouding Condition holds for all firms and zero profits otherwise. An increase in Δa_n for some $n\neq 1$ increases the left-hand-side of the Shrouding Condition and hence relaxes it; thus it either does not affect firm 1's profits or—if it makes the Shrouding Condition hold for some firm $n\neq 1$ for which it does not hold otherwise—strictly increases firm 1's profits. $\square$

Proof of Proposition 7. We solve for the equilibria of the subgames following firm 1's innovation decision. Absent innovation, the shrouding condition is satisfied as we are in a socially wasteful industry. Hence, a shrouded-prices equilibrium exists and (using our selection criterion that firms play this equilibrium in the pricing subgame whenever it exists) firm 1 therefore earns $s_1(\underline{f}+\overline{a}-c_1)$.

Now consider the subgame following a decision to innovate by firm 1. We first establish that there exists an equilibrium in which firm 1 offers the contract $(\underline{f}+[\Delta v_1-\max_{n\neq 1}\Delta v_n],\overline{a})$ with probability one, and all firms $n\neq 1$ offer the contract $(\underline{f},\overline{a})$; in this equilibrium all consumers are indifferent between firm 1 and its best competitor and following our tie-breaking rule buy firm 1's product. Since unshrouding yields zero profits in a socially wasteful industry, it is immediate that there exist no deviation for a firm $n\neq 1$ that yields positive profits. If firm 1 deviates and unshrouds or shrouds and sets a higher base price, it earns zero profits. And since it has a market share of one, firm 1 cannot benefit from lowering its base or additional price. Hence, firm 1 also plays a best response.

To complete the proof, we show that in any pricing subgame following innovation in which firms shroud with probability 1, firm 1 charges prices $\underline{f}+[\Delta v_1-\max_{n\neq 1}\Delta v_n],\overline{a}$ with probability one and gets the entire market, so that our equilibrium-selection criterion selects such an equilibrium.

This means that firm 1 earns $\underline{f} + [\Delta v_1 - \max_{n\neq 1} \Delta v_n] + \overline{a} - c_1$ if it innovates and $s_1(\underline{f} + \overline{a} - c_1)$ if it does not. I_v^* is the difference between these two profit levels.

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

Case I: Firm $\hat{n}$ sets $\overline{f}_{\hat{n}}$ with positive probability. Then it is not a best response for firm 1 to set an up-front price f_1 above $\overline{f}_{\hat{n}} + \Delta v_1 - \Delta v_{\hat{n}}$ because with such base prices firm 1 earns zero profits while it earns positive profits when offering a contract $(\underline{f}, \overline{a})$. Thus, firm 1 sets base prices $f_1 \leq \overline{f}_{\hat{n}} + \Delta v_1 - \Delta v_{\hat{n}}$, contradicting the fact that firm $\hat{n}$ wins with positive probability when setting $\overline{f}_{\hat{n}}$.

Case II: Firm $\hat{n}$ sets $\overline{f}_{\hat{n}}$ with zero probability. Hence, for every $\epsilon > 0$, firm $\hat{n}$ sets base prices in the interval $(\overline{f}_{\hat{n}} - \epsilon, \overline{f}_{\hat{n}})$ with positive probability; and this probability goes to zero as $\epsilon \to 0$. Let $\gamma \leq 1$ be the probability that all firms shroud. That firm $\hat{n}$ earns positive profits implies that $\gamma > 0$. Then, firm 1 earns equilibrium profits of at least $\gamma s_1(\underline{f} + \overline{a} - c_1) > 0$, which it can ensure by shrouding and offering the contract $(\underline{f}, \overline{a})$. Since as $\epsilon \to 0$ firm 1's profits go to zero when setting a base price at or above $\overline{f}_{\hat{n}} - \epsilon + \Delta v_1 - \Delta v_{\hat{n}}$, there exists an $\overline{\epsilon} > 0$ such that firm 1 earns lower profits when setting a base price at or above $\overline{f}_{\hat{n}} - \overline{\epsilon} + \Delta v_1 - \Delta v_{\hat{n}}$ than when shrouding and offering the contract $(\underline{f}, \overline{a})$. Hence, firm 1 sets base prices at or below $\overline{f}_{\hat{n}} - \overline{\epsilon} + \Delta v_1 - \Delta v_{\hat{n}}$, contradicting the fact that firm $\hat{n}$ wins with positive probability when setting prices in the interval $(\overline{f}_{\hat{n}} - \overline{\epsilon}, \overline{\epsilon})$.

Finally, we show that in any equilibrium in which all firms $n \neq 1$ shroud with probability 1, firm 1 shrouds and offers the contract $(\underline{f} + [\Delta v_1 - \max_{n\neq 1} \Delta v_n], \overline{a})$ with probability one; hence, consumers weakly prefer fir m 1, and firm 1 gets the entire market. If firms $n \neq 1$ shroud with positive probability, firm 1 can earn positive profits by shrouding and offering the above contract. Hence firm 1 shrouds with probability one. Furthermore, since firm 1 makes positive profits only conditional on all rivals shrouding, it must set $a_1 = \overline{a}$ in any such equilibrium. Since conditional on all firms shrouding, firm 1 attracts all consumers with probability 1 when setting $\underline{f} + [\Delta v_1 - \max_{n \neq 1} \Delta v_n]$, firm 1 does not charge a lower base price in such an equilibrium. Finally, firm 1 cannot charge strictly more than $\underline{f} + [\Delta v_1 - \max_{n \neq 1} \Delta v_n]$ with positive probability because otherwise some firm $n \neq 1$ could make positive profits when shrouding and offering the contract $(\underline{f}, \overline{a})$, which contradicts the fact that all firms $n \neq 1$ earn zero profits. □

Proof of Proposition 8. Absent innovation, firm 1 earns $s_1(\underline{f} + \overline{a} - c_1)$ in the shrouded-prices equilibrium we select. From now on, consider the subgame following firm 1's decision to innovate. First, observe that there exists an equilibrium in which all firms unshroud with probability one. In this case the game simplifies to a standard Bertrand game. Since we focus on equilibria in which no firm prices below marginal costs, firm 1 earns $RA_1(v + \Delta v_1, \ldots, v + \Delta v_N)$ in such an equilibrium.

We next argue that firm 1 earns at least $RA_1(v + \Delta v_1, \ldots, v + \Delta v_N)$ in any equilibrium. Recall that by assumption no firm prices below marginal cost. Thus, if firm 1 unshrouds and and charges a total price of $c_1 + RA_1(v + \Delta v_1, \ldots, v + \Delta v_N)$ all consumers weakly prefer the product of firm 1 and our tie-breaking assumption implies that firm 1 serves the entire market, earning its relative advantage.

We will now argue that firm 1 earns no more than its relative advantage. Suppose otherwise. Since firm 1 earns more than its relative advantage it must do so (in expectation) for all but a set of measure zero of total prices it charges; hence there exists an $\epsilon > 0$ such that firm 1 charges a total price above $c_1 + RA_1(v + \Delta v_1, \ldots, v + \Delta v_N) + \epsilon$ for some $\epsilon > 0$ with probability 1. Let k be a firm for which $\Delta v_k - c_k = \max_{n \neq 1}[\Delta v_n - c_n]$. Firm k must earn positive profits; otherwise it could deviate, unshroud and offer a total price of $c_k + \epsilon/2$, thereby offering a better deal to consumers than firm 1 and hence win with positive probability and earn positive profits. Furthermore, since the equilibrium outcome does not coincide with that of the corresponding standard Bertrand game, all firms must shroud with positive probability.

Let $\hat{t}_k$ be the supremum of the total price distribution firm k charges; and let $\hat{t}_1$ be that of firm 1. Define the quality adjusted maximum of these suprema as $\hat{t} = \max\{\hat{t}_1 - \Delta v_1, \hat{t}_k - \Delta v_k\}$. Note that firm k cannot charge this quality-adjusted total price with positive probability when unshrouding; if it did, it would lose to firm 1 with probability one, contradicting the fact that it must earn positive profits with any price it charges with positive probability. Furthermore, firm k cannot charge this quality-adjusted price with positive probability when shrouding. If it did, it must have positive market share and it can do so only if all other firms shroud. But then it must set $a_k = \overline{a}$ to maximize profits, and hence it offers a contract $(\hat{t}+\Delta v_k - \overline{a}, \overline{a})$ with positive probability. For this contract to have a positive market share, firm 1 must shroud and set base prices above $\hat{t} + \Delta v_1 - \overline{a}$ with positive probability, and for its quality-adjusted price to be below $\hat{t}$ firm 1 at the same time must set the additional price below $\overline{a} - \Delta v_1$. But this is not a best response: firm 1 could keep the total price distribution fixed but always charge the maximal additional price $\overline{a}$. Then its base price would always lie (weakly) below $\hat{t} + \Delta v_1 - \overline{a}$, and firm 1 could strictly increase its market share holding the total price distribution fixed. We conclude that firm k cannot charge a quality-adjusted price of $\hat{t}$ with positive probability. But then firm 1 cannot do so when unshrouding as it would have zero market share. And firm 1 cannot charge this price with positive probability when shrouding: if it did, firm k would have to charge base price weakly above $\hat{t}+\Delta v_k - \overline{a}$ when shrouding for firm 1 to have positive market share. But firm k could keep its total price distribution fixed and always charge the maximal additional price $\overline{a}$, thereby ensuring that its base price is below $\hat{t} + \Delta v_k - \overline{a}$, increasing its market share while holding the total price fixed. We conclude that neither firm charges the highest quality-adjusted price with positive probability.

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

The above statement is immediate if the firm in question unshrouds, so that for a sufficiently small $\epsilon > 0$ it must be that firms 1 and k almost always shroud when they set total quality-adjusted prices in the interval $(\hat{t} - \epsilon, \hat{t})$. Suppose, then, that firm 1 shrouds on this interval, but its market share does not approach zero. Then firm k must with positive probability shroud and set up-front prices at or above $\hat{t} + \Delta v_k - \overline{a}$ while setting total prices strictly below $\hat{t} + \Delta v_k$. Firm k in this case could keep its total price distribution fixed, and always charge the maximal additional price $\overline{a}$, increasing its market share for a set of total prices it charges with positive probability. The argument for why the market share of firm k must go to zero when shrouding and charging quality-adjusted total prices in $(\hat{t}-\epsilon, \hat{t})$ is analogous. We conclude that firm 1 earns its relative advantage in every equilibrium.

Thus, $I_v^* = RA_1(v+\Delta v_1, \ldots, v+\Delta v_N) - s_1(\underline{f}+\overline{a}-c_1)$. If $RA_1(v+\Delta v_1, \ldots, v+\Delta v_N) = 0$, trivially the first inequality in the proposition holds and $I_v^* < 0$. Thus, suppose $RA_1(v+\Delta v_1, \ldots, v+\Delta v_N) > 0$. Then the first inequality in the proposition holds because

$$\begin{aligned}
I_v^* &= \Delta v_1 - c_1 - \max_{n\neq 1}(\Delta v_n - c_n) - s_1(\underline{f}+\overline{a}-c_1) \qquad (6)\\
&\leq \Delta v_1 - c_1 - \max_{n\neq 1}\Delta v_n + \max_{n\neq 1} c_n - s_1(\underline{f}+\overline{a}-c_1)\\
&= \big[(1-s_1)(\underline{f}+\overline{a}-c_1)\big] + \left[\Delta v_1 - \max_{n\neq 1}\Delta v_n\right] - (\underline{f}+\overline{a}-\max_{n\neq 1} c_n)
\end{aligned}$$

and $\underline{f}+\overline{a} > c_n$ for all n. To establish that $I_v^* < 0$, we use the fact that $-c_1 - \max_{n\neq 1}(\Delta v_n - c_n) \leq -c_1 - \min_{n\neq 1}\Delta v_n + c_{min} \leq 0$ in Equality (6) above. □

Proof of Proposition 9. 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 can earn positive profits by unshrouding the additional prices and offering $(\underline{f}, \overline{a})$.

(iii): *The distributions of total prices are bounded from above.* Suppose firm n sets the total price $f_n + a_n > v + \overline{a}$ with positive probability in equilibrium. When the additional prices are shrouded, consumers never buy the product from firm n because this inequality implies $f_n > v$. 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 $(\hat{f}+\hat{a})_n$ be the supremum of the equilibrium total-price distribution of firm n *conditional* on firm n unshrouding;

set $(\hat{f}+\hat{a})_n = 0$ in case firm n does not unshroud. Let $\hat{f}+\hat{a} = \max_n \left\{(\hat{f}+\hat{a})_n\right\}$. Consider firm n that unshrouds and for whom $(\hat{f}+\hat{a})_n = (\hat{f}+\hat{a})$. Note that in any equilibrium in which some firm unshrouds with positive probability, $\hat{f}+\hat{a} > c_n$ by (ii). Also, by (iii), $\hat{f}+\hat{a}$ is bounded from above and hence well-defined.

First, suppose that firm n charges the total price $\hat{f}+\hat{a}$ with positive probability. If some other firm $n' \neq n$ also sets the total price $\hat{f}+\hat{a}$ 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 $\hat{f}+\hat{a}$ with positive probability. Because $\hat{f}+\hat{a}$ 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 $\hat{f}+\hat{a} > v$. When firm n shrouds and offers $(\underline{f}, \overline{a})$, however, its profits conditional on all other firms shrouding are at least $s_n(\underline{f}+\overline{a}-c_n)$. Thus, the equilibrium condition $s_n(\underline{f}+\overline{a}-c_n) \geq v - c_n$ implies that deviating by shrouding and offering $(\underline{f}, \overline{a})$ is profitable—a contradiction.

Second, suppose that firm n does not charge the total price $\hat{f}+\hat{a}$ with positive probability. Then, for any $\epsilon > 0$, firm n charges a total price in the interval $(\hat{f}+\hat{a}-\epsilon, \hat{f}+\hat{a})$ 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 $\hat{f}+\hat{a}$ 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 looses money in expectation relative to shrouding and offering $(\underline{f}, \overline{a})$. In addition, conditional on all other firms shrouding firm n earns less than 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* $(\underline{f}, \overline{a})$ *with probability one.* By (iii), all firms choose to shroud with probability one. Hence, in equilibrium all firms charge an additional price $a = \overline{a}$ with probability one. By the exact same argument as in Proposition 2, all firms offer a base fee $f = \underline{f}$. □

Proof of Proposition 10. 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 with 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 the exploitative equilibrium exists if Inequality (4) holds. Suppose that Inequality (4) does not hold for firm $\tilde{n}$ 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 by setting its total price above at $t - \epsilon$ with sufficiently small $\epsilon > 0$. Also, there is no firm which has 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 that range of prices 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 $t_{n'} = \underline{f} + \overline{a}$ to $f_{n'} = \underline{f}$ and its additional price slightly lower than $\hat{a}$—a contradiction. Therefore, there is no exploitative equilibrium when Inequality (4) does not hold. □