

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays in Behavioral Industrial Organization

Permalink

<https://escholarship.org/uc/item/53x8x14x>

Author

Murooka, Takeshi

Publication Date

2014

Peer reviewed|Thesis/dissertation

Essays in Behavioral Industrial Organization

by
Takeshi Murooka

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy
in
Economics
in the
Graduate Division
of the
University of California, Berkeley

Committee in charge:
Professor Botond Kőszegi, Co-Chair
Professor Matthew Rabin, Co-Chair
Professor Stefano Della Vigna
Professor Edward Augenblick

Spring 2014

© 2014
Takeshi Murooka
All Rights Reserved.

Abstract

Essays in Behavioral Industrial Organization

by

Takeshi Murooka

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Botond Kőszegi, Co-Chair

Professor Matthew Rabin, Co-Chair

This dissertation consists of three chapters on behavioral industrial organization.

The first chapter, titled “Deception under Competitive Intermediation,” investigates the incentives of intermediaries—such as mortgage brokers, financial advisors, or insurance salespeople—to educate consumers who misperceive the value of products. Two types of firms sell products through competing common-agent intermediaries and pay commissions for sales. One sells a transparent product, while the other sells a deceptive product that has a hidden fee, quality, or risk. Each intermediary chooses which product to offer and whether or not to educate consumers about the hidden attribute. Each consumer visits a fixed number of intermediaries and buys at most one item. When consumers correctly anticipate the hidden attribute, intermediaries reveal it and commissions are competed away. When consumers misperceive the hidden attribute, however, intermediaries employ deception if and only if the degree of misperception is large. If deception occurs, intermediaries earn high commissions despite competition. Furthermore, because consumers ultimately bear the cost of such commissions, consumer welfare is lower when intermediaries can educate consumers than when they cannot. Deception is less likely to occur when consumers visit more intermediaries before making their purchase decisions. Conditional on deception, however, visiting more intermediaries further raises the level of commissions because deceptive firms need to give each intermediary a higher commission to maintain the deception. Regulating commissions—analogueous to recent policies in the US mortgage industry as well as in the Australian and UK mutual-fund industries—can lead intermediaries to reveal all hidden attributes.

The second chapter, titled “Inferior Products and Profitable Deception” and co-authored

with Paul Heidhues and Botond Kőszegi, analyzes conditions facilitating profitable deception in a simple model of a competitive retail market. Firms selling homogenous products set up-front prices that consumers understand and additional prices that naive consumers ignore unless revealed to them by a firm, where—to model especially financial products such as credit cards and mutual funds—we assume that there is a binding floor on the up-front prices. Our main results establish that “bad” products (those with lower social surplus than an alternative) tend to be more reliably profitable than “good” products. Specifically, (1) in a market with a single socially valuable product and sufficiently many firms, a deceptive equilibrium—in which firms hide additional prices—does not exist and firms make zero profits. But perversely, (2) if the product is socially wasteful, then a firm cannot profitably sell a transparent product, so there is no incentive to reveal the additional prices and hence a profitable deceptive equilibrium always exists. Furthermore, (3) in a market with multiple products, since a superior product both diverts sophisticated consumers and renders an inferior product socially wasteful in comparison, it guarantees that firms can profitably sell the inferior product by deceiving consumers.

The third chapter, titled “Exploitative Innovation” and co-authored with Paul Heidhues and Botond Kőszegi, studies innovation incentives in a simple model of a competitive retail market with naive consumers. Firms selling perfect substitutes play a game consisting of an innovation stage and a pricing stage. At the pricing stage, firms simultaneously set a transparent “up-front price” and an “additional price,” and decide whether to shroud the additional price from naive consumers. To capture especially financial products such as banking services, credit cards, and mutual funds, we allow for a floor on the product's up-front price. At the preceding innovation stage, a firm can invest either in increasing the product's value (value-increasing innovation) or in increasing the maximum additional price (exploitative innovation). We show that if the price floor is not binding, the incentive for either kind of innovation equal the “appropriable part” of the innovation, implying similar incentives for exploitative and value-increasing innovations. If the price floor is binding, however, innovation incentives are often stronger for exploitative than for value-increasing innovations. Because learning ways to charge higher additional prices increases the profits from shrouding and thereby lowers the motive to unshroud, a firm may have strong incentives to make appropriable exploitative innovations, and even stronger incentives to make non-appropriable exploitative innovations. In contrast, the incentive to make non-appropriable value-increasing innovations is zero or negative, and even the

incentive to make appropriable value-increasing innovations is strong only if the product is socially wasteful. These results help explain why firms in the financial industry have been willing to make innovations others could easily copy, and why these innovations often seem to have included exploitative features.

Contents

1	Deception under Competitive Intermediation	1
1.1	Introduction	1
1.2	Related Theoretical Literature	5
1.3	Model	6
1.3.1	Setup	6
1.3.2	Discussion of Key Assumptions	8
1.4	Analysis	10
1.4.1	Benchmark Cases	10
1.4.2	Equilibria in the Model	11
1.4.3	Welfare Effects of Intermediaries	15
1.5	Policy Analysis	16
1.5.1	Regulating Commissions	17
1.5.2	Regulations on Hidden Attributes	19
1.5.3	Enhancing Access to Intermediaries	19
1.5.4	Mandatory Disclosure of Commission Structure	20
1.6	Extensions	20
1.6.1	Competition among Deceptive Firms	21
1.6.2	Heterogenous Consumers	22
1.6.3	Further Extensions and Modifications	25
1.7	Concluding Remarks	29
1.8	Proofs and Supplementary Materials	30
1.8.1	Proofs	30
1.8.2	Further Benchmark Cases	34
1.8.3	Costly Education	36
2	Inferior Products and Profitable Deception (with Paul Heidhues and Botond Kőszegi)	39
2.1	Introduction	39
2.2	Basic Model	42
2.2.1	Setup	42
2.2.2	Applications	43
2.2.3	Microfoundations for the Floor on the Up-Front Price	47
2.3	Profitable Deception	50

2.3.1	Benchmarks: Non-Binding Price Floor or Sophisticated Consumers	50
2.3.2	Naive Consumers with a Binding Price Floor	51
2.3.3	Socially Valuable versus Socially Wasteful Products	53
2.4	Sophisticated Consumers and Multi-Product Markets	54
2.4.1	Sophisticated Consumers in Our Basic Model	54
2.4.2	Sophisticated Consumers with an Alternative Transparent Product	55
2.5	Extensions and Modifications	59
2.5.1	Costly Unshrouding and Market Power	59
2.5.2	Inefficiencies in Collecting Revenue from Additional Prices	62
2.5.3	Misprediction of Value	63
2.5.4	Further Extensions and Modifications	64
2.6	Related Theoretical Literature	65
2.7	Some Policy Implications and Conclusion	67
2.8	Proofs	69
3	Exploitative Innovation (with Paul Heidhues and Botond Köszegi)	83
3.1	Introduction	83
3.2	Basic Model	86
3.3	Non-Binding Price Floor	88
3.4	Binding Price Floor	89
3.5	Concluding Remarks	93
3.6	Proofs and Supplementary Materials	94
3.6.1	Proofs	94
3.6.2	Uniqueness of Deceptive Equilibrium with Unshrouding Cost	98
	Bibliography	101

Acknowledgements

In preparing this dissertation, I have received enormous supports from my advisors, Botond Kőszegi and Matthew Rabin. I have learned countless things from Botond, and I cannot thank too much for his heartfelt mentorship. Matthew has advised me with his profound insights, and I am hugely indebted to his sincere guidance. I am also truly grateful to Stefano DellaVigna and Paul Heidhues for helpful suggestions, solid discussions, and warm encouragements.

I have also learned a lot from other faculty members, seminar and conference audiences at many places, and my (past and present) colleagues. Among my wonderful colleagues at Berkeley, I would especially like to thank Sheng Li, Aniko Öry, and Antonio Rosato. Financial supports during my doctoral study from the department, the Murata Overseas Scholarship Foundation, and the Program in Psychological Economics are cordially acknowledged.

I dedicate this dissertation to my parents, sisters, brothers, and wife.

Chapter 1

Deception under Competitive Intermediation

1.1 Introduction

In the mortgage, mutual fund, and insurance industries, products are often sold through independent intermediaries.¹ A primary role of the intermediaries is to help consumers make better purchase decisions by informing them about product attributes. This educational role of intermediaries is particularly important for uninformed or confused consumers who may be inattentive to “hidden” fees, qualities, or risks.² Nevertheless, recent empirical studies report that intermediaries often give advice that is detrimental to consumers but benefits product providers.³ Some of these studies find that intermediaries receive higher commissions from product providers for selling products which are worse for consumers.⁴

¹In the US, the Investment Company Institute reports that among all households who hold mutual funds including pension plans, 53 percent of them own funds purchased through investment professionals, and 82 percent of households do so after excluding pension plans (Profile of Mutual Fund Shareholders, 2012). The Mortgage Bankers Association (MBA) reports that 50 percent of all mortgage loans and 71 percent of subprime loans are originated through mortgage brokers (Residential Mortgage Origination Channels, MBA Research Data Notes, 2006).

²Anagol and Kim (2012) find that investors are less sensitive to mutual-fund fees when the fees are amortized and hidden. Gurun, Matvos and Seru (2013) report that consumers are less sensitive to post-introductory interest rates than to initial interest rates of adjustable-rate mortgages because of “deceptive advertisements” by mortgage lenders. See also the Federal Trade Commission’s article on deceptive mortgage advertisements: <http://www.consumer.ftc.gov/articles/0087-deceptive-mortgage-ads>.

³Mullainathan, Nöth and Schoar (2010) conduct an audit study and find that most financial advisors cater to their customers’ biases, such as return chasing, and promote high-fee mutual funds. Anagol, Cole and Sarkar (2012) report that 60 percent or more of life-insurance salespeople recommend strictly-dominated insurance plans.

⁴Chalmers and Reuter (2012) report that customers who consulted brokers for retirement plans allocate their money more to funds with higher broker fees, although on average these broker-recommended funds underperform a default investment option. Christoffersen, Evans and Musto (2013) find that in the US mutual-fund industry, a higher commission increases a fund flow while it also predicts future poorer fund performance.

Yet, how intermediaries can profitably sell worse products *and* get higher commissions in a competitive environment remains largely unexplored.

Building on Gabaix and Laibson (2006) and complementing the literature on intermediation under consumer naivete (Stoughton, Wu and Zechner 2011; Bolton, Freixas and Shapiro 2012; Inderst and Ottaviani 2012c), this paper investigates the incentives of competing intermediaries to educate consumers who misperceive the value of products. I show that when intermediaries are motivated by commissions, deception (i.e., not educating consumers about their misperception) occurs if and only if the degree of misperception is large. In the deceptive equilibrium, each intermediary faces a trade-off between expanding market share by educating consumers and earning higher commissions per sale by exploiting consumers. Based on this trade-off, intermediaries engage in deception if deceptive firms can pay sufficiently high commissions—financed by deception—with which transparent firms cannot compete. If deception occurs, then intermediaries receive high commissions even when they are competing for consumers. Such deception severely harms social and consumer welfare. Consistent with the evidence described in the previous paragraph, intermediaries can profitably sell products with lower, or even negative, social surplus. Intermediaries are less likely to educate consumers when their educational role is more important. Further, consumer welfare is lower when intermediaries can—but do not—educate consumers than when they cannot educate consumers. Analogous to recent policies in the US mortgage industry as well as in the Australian and UK mutual-fund industries, regulating commissions can lead intermediaries to educate consumers and hence can increase consumer and social welfare.⁵

After summarizing the related theoretical literature in Section 1.2, I set up the model and discuss its key assumptions in Section 1.3. In the model, two firms sell their products to a unit mass of homogenous consumers. One firm produces a deceptive product that has a hidden product attribute such as an additional fee, a harmful quality, or a future risk, whereas the other firm produces a transparent product that has no hidden attribute. Firms can sell their products only through profit-maximizing common-agent intermediaries, to whom they pay sales commissions. Each intermediary decides which product to promote and whether to educate consumers about the hidden attribute of the deceptive product. Each consumer visits a fixed number of intermediaries and buys at most one item. Following Gabaix and Laibson (2006) and Heidhues, Köszegi and Murooka (2012a), I assume that consumers are naive both in the sense that they are initially unaware of the hidden attribute and that they do not infer its existence from the level of prices or commissions. Consumers take hidden attributes into account when making their purchase decisions if and only if they are educated by some intermediary. I investigate subgame-perfect Nash equilibria played by firms and intermediaries. In particular, I focus on identifying conditions for equilibria in which intermediaries employ deception.

Section 1.4 analyzes the model and discusses welfare implications. To compare equilib-

⁵In the US mortgage industry, “to protect mortgage borrowers from unfair, abusive, or deceptive lending practices,” the Federal Reserve Board has prohibited compensation to a mortgage broker based on terms or conditions of a mortgage transaction since 2011 (Banking and Consumer Regulatory Policy Press Release, August 16, 2011). Also, in the Australian and UK mutual-fund industries, commissions to financial advisors have been banned since 2013.

rium outcomes across market structures, before the main analysis I investigate two benchmark cases. First, I show that if firms can directly educate consumers about hidden attributes, then the transparent firm always reveals the deceptive firm's hidden attributes. Second, I show that if consumers are sophisticated in that they either know which products have hidden attributes or anticipate that some products have hidden attributes, then a product with lower social surplus is never sold and commissions are driven down to a competitive level.

I then investigate the main model in which consumers are naive and firms sell their products through intermediaries. Holding the other parameters constant, I show that deception occurs if and only if the amount of the hidden attribute is large. Specifically, the condition for deception hinges on an intermediary's trade-off between expanding market share and earning higher commissions per consumer. On the one hand, an intermediary can increase its market share by educating consumers and attracting them from other intermediaries. On the other hand, an intermediary can earn higher commissions by not educating consumers and selling the deceptive product. As a result, deception occurs if the deceptive firm can give sufficiently high commissions—financed by the hidden attribute—with which the transparent firm cannot compete. Because the deceptive firm needs to give a sufficiently high commission to each intermediary to maintain the deception, competition among intermediaries does not lower the level of commissions when deception occurs. Conditional on deception, increasing consumers' search intensity (i.e., the number of intermediaries consumers visit) further raises the level of commissions. Nevertheless, deception becomes less likely to occur as the search intensity increases. Hence, consumers' search intensity has a non-monotonic effect on the level of commissions.

When deception occurs, the educational role of intermediaries exhibits perverse welfare effects. Deception distorts consumer and social welfare; because consumers misperceive the value of products, intermediaries can profitably sell products with lower social surplus or even ones with negative social surplus. Intermediaries are less likely to educate consumers as their educational role becomes more important (i.e., as the hidden attribute is larger). Further, I show that consumer welfare is lower when intermediaries can educate consumers about the hidden attribute than when they cannot. This is because commissions for persuading intermediaries not to educate consumers increase the total prices of the products, and consumers ultimately bear the cost of such commissions. This result indicates that if deception is an issue, having expert intermediaries in a market can hurt naive consumers. I also show that conditional on deception, the ex-post utility of consumers is the same under a monopoly intermediary and multiple intermediaries. Although introducing competition among intermediaries makes deception harder to maintain, it does not increase consumer or social welfare if deception is maintained.

Section 1.5 discusses the possibilities and limits of policies for preventing deception. Once the difference in commissions is limited, intermediaries would try to attract consumers from competitors by educating the consumers. Therefore, caps on commissions or prohibiting large discrepancies in commissions can eliminate deception, and thereby increase welfare. This is akin to recent regulations introduced in the US mortgage industry as well as in

the Australian and UK mutual-fund industries. Unlike policies that attempt to restrict hidden attributes directly, these commission regulations do not require a policymaker to identify which attribute is used to exploit consumers.⁶ I also discuss the effects of regulating the maximum additional fees, letting consumers reach more intermediaries, and whether commission structures of intermediaries are disclosed or not.

Section 1.6 investigates two major extensions and discusses other modifications of the model. As one major extension, I analyze how competition among many firms affects consumer and social welfare. When there are multiple firms in each type of product, all firms earn zero profits. In this case, whether or not intermediaries earn positive profits from deception depends on the relative social surplus of the products. On the one hand, when the deceptive product is socially superior to the transparent one, deceptive firms compete down their product prices and commissions, intermediaries make zero profits, consumers' ex-post utility is positive, and social welfare is maximized. On the other hand, when the deceptive product is socially inferior to the transparent product—which seems more likely in practice—intermediaries can earn positive profits by employing deception. The same trade-off and condition as in the model with one deceptive firm determine whether deception through high commissions can be maintained. If the incentive for receiving high commissions outweighs the incentive for expanding market share, then intermediaries earn positive profits from deception. Each consumer's ex-post utility is higher than in the case with one deceptive firm, but is still negative.

As another major extension, I investigate how the presence of sophisticated consumers affects the welfare of naive consumers. Suppose that a fraction of consumers are *informed* in the sense that they know the existence of hidden attributes and which products have the hidden attributes. I first show that deception through high commissions can still occur if each intermediary is able to offer only one product at a time. Nevertheless, the condition for deception becomes more stringent as the fraction of informed consumers increases. I next examine a model in which each intermediary can screen consumers by offering a menu that contains one promoted product and other non-promoted products. I show that if intermediaries can conceal both the existence of non-promoted products as well as the hidden attributes of deceptive products from naive consumers, then intermediaries can profitably sort consumers. Naive consumers buy inferior deceptive products with high commissions which are promoted by intermediaries, whereas informed consumers buy superior transparent products with low commissions which are available but are not advertised by the intermediaries.

I then discuss further extensions and modifications of the model incorporating: (i) effort costs intermediaries need to pay when educating consumers, (ii) heterogeneity in consumers' search intensity, (iii) the possibility of vertical integration, (iv) the possibility that intermediaries directly charge advising fees or give direct rebates to consumers, and (v) the possibility

⁶Furthermore, this policy has a positive effect on a relevant issue: preventing firms from inventing new hidden attributes of which policymakers are not aware. If commissions are regulated, then intermediaries would detect and reveal all hidden attributes. This would eliminate firms' incentives to invent new hidden attributes. Hence, intermediaries can work as monitoring institutions once commissions do not distort their incentives.

of promoting multiple products. Section 1.7 concludes. The proofs are provided in the Appendix.

1.2 Related Theoretical Literature

This section summarizes related theoretical literatures. I first discuss studies that are most closely related to this paper: markets with intermediation under consumer naivete. I also review studies that investigate competition under consumer naivete in a retail market, and that analyze the role of intermediaries as information providers under rational consumers. Compared to these literatures, I show that intermediaries can earn high commissions by employing deception even when they compete for consumers, identify perverse welfare effects of intermediaries who are able to educate naive consumers, and shed light on new positive aspects of commission regulations.

This paper is most closely related to a growing literature that analyzes markets with intermediaries under consumer naivete. Stoughton, Wu and Zechner (2011) investigate a model with a monopolistic financial intermediary and show that commissions are used either for price discrimination across individual wealth levels or for socially-inefficient marketing, depending on the degree of investor naivete. Bolton, Freixas and Shapiro (2012) analyze competition among credit-rating agencies with trusting investors who always take the ratings at face value. Because firms may disclose only the most favorable rating to attract naive investors, the presence of multiple (truth-telling) credit-rating agencies facilitates ratings shopping of the firms and distorts social welfare. Hence, social welfare under duopoly credit-rating agencies can be lower than that under a monopoly credit-rating agency. Inderst and Ottaviani (2012c) analyze a market with a monopolistic intermediary and horizontally-differentiated firms. The authors show that when consumers are naive, the intermediary charges no advising fees to consumers directly but earns high commissions provided from firms, which leads to biased advice to the consumers.

This paper, as well as the papers summarized in the previous paragraph, builds on the theoretical literature investigating the effects of consumer naivete.⁷ Specifically, this paper assumes that consumers have misperceptions about certain product attributes but experts can educate them. Gabaix and Laibson (2006) develop a model with “educable” consumer naivete in a retail market. They analyze a model in which each firm sells a base product and an add-on. Naive consumers are initially inattentive to the prices of add-ons, but each firm can choose whether or not to inform both sophisticated and naive consumers about all prices of the add-ons. Because sophisticated and informed-naive consumers can substitute away from the add-on before buying a base product, such information disclosure

⁷See, for instance, DellaVigna and Malmendier (2004), Eliaz and Spiegel (2006, 2008), Spiegel (2006a, 2006b), Grubb (2009, 2012), Heidhues and Köszegi (2010), Piccione and Spiegel (2012), Inderst and Ottaviani (2013), and Gabaix, Laibson, Li, Li, Resnick and de Vries (2012). Relatedly, Gennaioli, Shleifer and Vishny (2012) analyze a retail fund market where investors perceive the variance of a risky asset is smaller as they trust a fund manager more. The distribution of trust makes the managers horizontally-differentiated and enables them to charge a fee above their marginal cost. The trusted managers tend to charge higher fees for riskier assets.

can decrease the demand for the add-on and hence may not be profitable for firms even under competition.⁸ Building upon this insight, Heidhues et al. (2012a) investigate retail markets with a floor on a base-product price, analyze a screening problem between sophisticated and naive consumers by offering multiple products, and identify the role of socially-inferior products for maintaining profitable deception.

This paper also belongs to the literature that analyzes the role of intermediaries as information providers.⁹ Lizzeri (1999) investigates an information-disclosure problem under a monopolistic intermediary. He also shows that competition among intermediaries can lead to full information disclosure. Inderst and Ottaviani (2009) analyze how the quality of advice can be distorted from the socially optimal level when a monopolistic intermediary pays a private cost to find a potential customer. Inderst and Ottaviani (2012a) investigate a market with a monopolistic intermediary and horizontally-differentiated product providers. The authors show that the mandatory disclosure of commission levels can distort the efficient provision of the products when there is a cost asymmetry between firms. This is because the market share of a cost-efficient firm is below the social optimum before commission disclosure, and the disclosure further reduces the equilibrium product supply of the cost-efficient firm.

1.3 Model

This section introduces the model. Section 1.3.1 sets up the model. Section 1.3.2 discusses three key assumptions of the model throughout this paper.

1.3.1 Setup

Consider a market with two product providers: a deceptive firm (firm D) and a transparent firm (firm T). Firm D charges a hidden fee $\bar{a} \geq 0$, whereas firm T charges no hidden fees. Firm $x \in \{D, T\}$ sells product x with value $v_x > 0$ and marginal cost $c_x > 0$. Assume that $v_D - c_D + \bar{a} > 0$ and $v_T - c_T > 0$.¹⁰ There is a unit mass of homogenous consumers; each of them buys at most one item. Consumers are *naive* but *educable* as in Gabaix and Laibson (2006) and Heidhues et al. (2012a): when consumers make purchase decisions, they are ignorant of \bar{a} if and only if they are not educated about \bar{a} . For simplicity, I assume that \bar{a} is exogenous and consumers cannot avoid it after the purchase.¹¹ Note that firm D has monopoly power for potentially exploiting consumers by \bar{a} ; Section 1.6.1 analyzes a model with multiple deceptive firms in which no firm has any monopoly power.

⁸See also Miao (2010), Armstrong and Vickers (2012), Li, Peitz and Zhao (2012), Dahremöller (2013), Warren and Wood (2010), Kosfeld and Schüwer (2011), and Ko and Williams (2011).

⁹See Gorton and Winton (2003), Dranove and Jin (2010), and Inderst and Ottaviani (2012b) for review.

¹⁰Otherwise, some product is never profitably sold and the market becomes a monopoly. Note that product D can be socially wasteful (v_D can be smaller than c_D). Note also that I do not impose a specific relation between the social surplus of these two products ($v_D - c_D$ versus $v_T - c_T$).

¹¹If instead the hidden fee is avoidable and endogenously chosen by firm D , then the firm sets the hidden fee equal to a monopoly price after consumers are locked-in.

A key feature of the model is that firms must delegate their sales to common-agent intermediaries. Let $J \geq 2$ denote the total number of intermediaries in the market. Each consumer randomly visits a fixed number $N(\leq J)$ of intermediaries.¹² I assume $N \geq 2$ to analyze a competitive environment of intermediaries; N limits each intermediary's ability to take market share away from competitors.¹³ Each intermediary chooses one product to promote, and whether or not to educate consumers about \bar{a} . Each intermediary can educate all consumers who visit at zero cost. If no intermediary educates, then a consumer is ignorant of \bar{a} in her purchase decision; but if at least one intermediary educates, then she takes \bar{a} into account. I assume that consumers do not make an inference about the hidden attribute from the level of product prices or commissions.¹⁴ All parties are risk neutral. I employ a tie-breaking rule where intermediaries split the demand equally if they promote the same product and consumers are indifferent between buying from them.¹⁵

The timing of the game is as follows:

1. Each firm $x \in \{D, T\}$ simultaneously proposes a product price p_{xi} and a commission per sale f_{xi} to each intermediary $i \in \{1, \dots, J\}$.¹⁶
2. After observing all of the contracts, each intermediary simultaneously chooses one product to promote, and whether or not to educate consumers about \bar{a} .
3. Each consumer reaches N intermediaries simultaneously and randomly, and makes a purchase decision.
4. All transactions are implemented.

The profits of firm D and T are respectively $p_{Di} - c_D - f_{Di} + \bar{a}$ and $p_{Ti} - c_T - f_{Ti}$ per sale. The ex-post utility of each consumer if she buys product D and product T from intermediary i is respectively $v_D - p_{Di} - \bar{a}$ and $v_T - p_{Ti}$. The total profits of each intermediary are its market share times commissions. It is worth mentioning that the model captures a case in

¹²In Section 1.6.3, I examine a model with incorporating heterogeneity in consumers' search intensity. Throughout this paper, I assume that the number of intermediaries each consumer visits is exogenous to analyze the effects of competition among intermediaries in a tractable way. Incorporating endogenous consumer search into the model is beyond the scope of this paper, though it is briefly discussed in Section 1.7.

¹³According to a survey reported by Lacko and Pappalardo (2007), in the US mortgage industry, consumers on average contact 2.8 mortgage lenders and brokers. Also, Woodward and Hall (2012) estimate that most consumers are likely to visit only 2 mortgage brokers for their loan originations.

¹⁴Incorporating commission-disclosure decisions into the model does not change the analysis. For ease of exposition, I consider a case in which consumers can observe the level of commissions but do not make an inference from it. See Section 1.5.4 for a detailed discussion.

¹⁵Further extensions are investigated in later sections. In Section 1.6.2, I incorporate consumer heterogeneity in naivete and analyze how intermediaries can screen the consumers' degree of naivete by a menu offer. In Section 1.6.3, I discuss a model that incorporates positive costs for educating consumers, vertical integration, advising fees or perks intermediaries can directly set to consumers, and promoting multiple products.

¹⁶For ease of exposition, I restrict the attention to piece-rate contracts. Given the demand structure, this restriction is without loss of generality in the model.

which \bar{a} is an overestimate of quality or underestimate of risk instead of a hidden fee. In such a case, a deceptive firm charges a high product price instead of an additional fee. Specifically, consider an alternative case where uneducated consumers perceive the valuation of product D to be $v_D + \bar{a}$, whereas its actual valuation is v_D . Then, all results from the original model remain the same once the product price of the deceptive firm is modified from p_{Di} to $p_{Di} + \bar{a}$.

I investigate pure-strategy subgame-perfect Nash equilibria played by firms and intermediaries with the following two equilibrium refinements. First, no firm sets its total price below its total cost. Indeed, any such strategy is weakly dominated. Second, in any off-equilibrium subgame, each intermediary takes the same educational action as it takes on the equilibrium path if it is a best response. This rules out off-equilibrium coordinations of educational decisions among intermediaries.¹⁷

For ease of exposition, I divide the set of equilibria into two types: *deceptive equilibria* in which some consumers remain uneducated, and *non-deceptive equilibria* in which all consumers are educated. In the analysis, I particularly focus on identifying conditions for and properties of deceptive equilibria. Since educating consumers is trivially a best response if all other intermediaries educate, a non-deceptive equilibrium always exists. Whenever a deceptive equilibrium exists, however, it is more plausible to be played between intermediaries than the non-deceptive equilibrium because of the following reasons. First, intermediaries earn higher profits in a deceptive equilibrium. Second, intermediaries play a weakly-dominated strategy in a non-deceptive equilibrium. Finally, if a deceptive equilibrium exists in the model, then it becomes the *unique* equilibrium in an extended model in which intermediaries incur a positive education cost, no matter how small the education cost is. I discuss such an extended model in Section 1.6.3.

1.3.2 Discussion of Key Assumptions

The model has three key assumptions: (i) consumers have misperceptions about a product attribute, (ii) intermediaries can educate consumers about the attribute, and (iii) without the help of intermediaries, firms cannot educate consumers about the attribute of other firms' products. In this subsection, I discuss these assumptions in turn.

(i) In the model, \bar{a} represents the amount by which a consumer misperceives the attributes of the product that can be hidden fees, qualities, or risks. As examples of hidden fees, Gurun et al. (2013) report that post-introductory interest rates of adjustable-rate mortgages are not salient due to “deceptive advertisements,” and the advertisements lead consumers to choose worse mortgages. Woodward and Hall (2012) find that some consumers originating mortgage loans pay high broker fees because of a confusing payment scheme.¹⁸ By examining

¹⁷In other words, given other intermediaries' strategies, if an intermediary chooses to educate (resp. not educate) on the equilibrium path, then the intermediary keeps choosing to educate (resp. not educate) in any other subgame whenever doing so is optimal. This refinement is akin to “no off-equilibrium signaling” assumption employed in incomplete-information games.

¹⁸Specifically, Woodward and Hall (2012) report that consumers who compensate a mortgage broker with both a direct cash payment and a commission from a mortgage lender pay twice as much as similar consumers who pay with cash alone or with a commission alone.

a natural experiment in the Indian mutual-fund industry, Anagol and Kim (2012) show that consumers tend to pay higher fees to mutual funds when the fees are amortized and hidden. As examples of misperceived qualities and risks, individual investors may overestimate future returns or underestimate risks of actively-managed mutual funds relative to index funds.¹⁹ Consumers may have incorrect beliefs about the likelihood of accidents covered by insurance plans. Patients may think the efficacy of a brand-name drug is better than that of generic one with exactly the same ingredients.

Along with most studies incorporating consumer naivete, I assume that consumers do not make an inference about the hidden attribute from price or commission levels. Of course, if consumers can rationally infer, then they will notice the existence of hidden attributes when observing overly high commissions. Empirical evidence, however, suggests that consumers are often inattentive to the incentives of intermediaries.²⁰ I return to discuss this assumption and policies on mandatory disclosure of commission structures in Section 1.5.4.

(ii) Helping consumers choose products is thought to be a central role of intermediaries. Doctors can teach patients which treatment is better for them, real-estate agents can tell deficiencies of a house, and financial advisors and mortgage brokers can educate consumers about the hidden costs of products. Experts in these industries are often indispensable because most consumers find it hard to choose appropriate products without the help of intermediaries. In addition, these intermediaries can provide certified information or clear analysis to modify consumers' misperceptions, whereas providing such information is either costly or often impossible for non-experts.

To investigate the educational incentives of intermediaries in a clear manner, I assume that each intermediary can educate its customers at no cost. Note that such "perfect education" is an extreme assumption which makes a deceptive equilibrium harder to exist. In Section 1.6.3, I examine how results are robust to incorporating costly education.

(iii) This paper focuses on markets in which expert intermediaries are indispensable for some consumers. Section 1.4.1 demonstrates that if firms can directly educate most consumers about other firms' product attributes, then a non-deceptive firm would always educate. For the industries illustrated above, however, some consumers are unwilling to buy products without consulting experts because stakes are large and product attributes are complicated. For example, mortgages have hundreds of thousands of dollars at stake, and their contracts are hundreds of pages long—far beyond the limits of comprehension for many consumers. To educate consumers in these markets, a non-deceptive firm needs either to hire or train in-house intermediaries. In either case, the total cost seems the same as, or higher than, that of using existent intermediaries. In Section 1.6.3, I discuss how results of

¹⁹Studies by Malkiel (1995), Gruber (1996), and French (2008) report that actively-managed mutual funds underperform index funds after fees are taken into account. Furthermore, Gil-Bazo and Ruiz-Verdú (2009) report that mutual funds charging higher fees tend to have worse before-fee risk-adjusted performance.

²⁰Malmendier and Shanthikumar (2007) report that small investors literally follow the stock recommendations of security analysts, though the recommendations of the analysts have an upward bias. Christoffersen et al. (2013) report that in the US mutual-fund industry, a 1% point increase in commissions leads to a 0.4464% increase in annual fund flows, while the increase in commissions predicts a 0.34% *decrease* in future performance net of fees.

the model are robust to incorporating such possibilities of vertical integration.²¹

1.4 Analysis

This section analyzes equilibria of the model and derives welfare implications. Section 1.4.1 presents two benchmark cases. Section 1.4.2 investigates the model, identifies a condition under which a deceptive equilibrium exists, and discusses its implications. Section 1.4.3 analyzes welfare effects on intermediaries' educational role and on the presence of competition among intermediaries.

1.4.1 Benchmark Cases

Before the main analysis, I first analyze two benchmark cases: a case where firms can directly market to consumers and a case where consumers do not have misperceptions. In Section 1.4.3, I investigate further benchmark cases—a case where neither firms nor intermediaries can educate consumers and a case where intermediaries have monopoly power—and discuss their welfare implications.

Equilibrium When Intermediaries Are Not Necessary

I first analyze a benchmark case in which consumers are naive and firms directly market to the consumers, which is a variant of an extended model in Heidhues et al. (2012a) where some firm produces only a socially-superior transparent product. Firm $x \in \{D, T\}$ simultaneously chooses its price p_x and whether or not to educate consumers about the hidden attribute of firm D .²² In this case, there always exists a Nash equilibrium played by firms in which firm T educates consumers about \bar{a} and a firm with a lower social surplus chooses marginal-cost pricing. This comes from the fact that once consumers are educated, the game reduces to one with Bertrand-type price competition in a vertically-differentiated market.²³

Suppose there exists a deceptive equilibrium. Since firms are facing Bertrand-type price competition, in equilibrium each consumer is indifferent between buying product D and T without taking \bar{a} into account:

$$v_D - p_D^* = v_T - p_T^*.$$

²¹Beyond the model, it is possible that non-deceptive firms can use mass advertisements to educate consumers. In that case, however, deceptive firms and intermediaries can also use “counter-advertisements.” Further, if profitable deception can occur, then deceptive firms and intermediaries have more resources to make naive consumers confused. Hence, education can be difficult without a direct consultation with an expert.

²²In the Supplementary Material, I show that how the result of Proposition 1.1 is robust to the different specifications of timing between pricing and educating decisions.

²³Precisely, there exists a non-deceptive equilibrium such that $p_D^* + \bar{a} = c_D, p_T^* = \min\{v_T, v_T - (v_D - c_D)\}$ and all consumers buy firm T 's product if $v_D - c_D \leq v_T - c_T$, whereas $p_D^* + \bar{a} = v_D - (v_T - c_T), p_T^* = c_T$ and all consumers buy firm D 's product if $v_D - c_D > v_T - c_T$.

In addition, at least one firm employs marginal-cost pricing; otherwise, some firm would profitably undercut the other firm. That is, either $p_D^* = c_D - \bar{a}$, $p_T^* = v_T - (v_D - c_D) - \bar{a}$ or $p_D^* = v_D - (v_T - c_T)$, $p_T^* = c_T$ holds in any equilibrium. In either case, by educating consumers, firm T can charge a higher price and still attract all of them. Proposition 1.1 summarizes the result:

Proposition 1.1 (Equilibrium When Intermediaries Are Not Necessary). Suppose firms directly market to consumers, and make pricing and educating decisions at the same time. Then, all consumers are educated in any equilibrium.

Equilibrium without Naivete

Next, suppose that firms sell their products through intermediaries but all consumers are *informed* about the hidden attribute. These informed consumers observe which product has \bar{a} . In this case, a standard Bertrand-type competition argument applies as summarized in Proposition 1.2.²⁴

Proposition 1.2 (Equilibrium without Naivete). Suppose firms market through intermediaries and all consumers are informed. Then, in any equilibrium, only the product with higher social surplus is sold and all intermediaries earn zero profits. The consumers' ex-post utility is non-negative.

In the Supplementary Material, I show that the result remains the same if instead consumers anticipate the existence of a hidden attribute but do not know which product has the hidden attribute. In this case, a firm who has no hidden attribute always induces intermediaries to educate consumers about the other firm's \bar{a} because the education would increase consumers' willingness to pay for its product. It leads that when intermediaries do not educate, they always promote a product with the hidden attribute. Hence, anticipated consumers always know which product has the hidden attribute based on the intermediaries' educational decisions.

1.4.2 Equilibria in the Model

Now I investigate the model presented in Section 1.3.1: firms sell products through intermediaries and consumers are naive. I first derive conditions for equilibria in which all intermediaries promote the deceptive product without educating consumers, and show that such a fully deceptive equilibrium is unique and is characterized by $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, $p_{Di}^* = v_D$, $f_{Di}^* = N(v_T - c_T)$.

Suppose an equilibrium exists in which all intermediaries promote the deceptive product without educating consumers. Let l be an intermediary to whom firm T proposes its lowest

²⁴Note that the proposition is stated in terms of utility and profits rather than what intermediaries actually do. There is a non-essential multiplicity of equilibria due to the fact that intermediaries make zero profits. This multiplicity affects none of the equilibrium outcomes.

product price: $p_{Tl} \equiv \min_s p_{Ts}$. First, $p_{Tl}^* \leq v_T$ in any equilibrium.²⁵ Also, $v_D - p_{Di}^* \leq v_T - p_{Tl}^*$ for all i ; otherwise firm D can profitably increase p_{Di}^* by a bit. These two inequalities imply that if intermediary l educates consumers and promotes product T , then all consumers who visit intermediary l strictly prefer to buy it from intermediary l . It leads that $(1/J)f_{Dl}^* \geq (N/J)f_{Tl}^*$; otherwise intermediary l would educate consumers and sell product T . Also, $(1/J)f_{Dl}^* \leq (N/J)f_{Tl}^*$; otherwise firm D can profitably decrease its commissions without inducing intermediaries' deviations. Thus,

$$f_{Dl}^* = Nf_{Tl}^*. \quad (1.1)$$

Given Equality (1.1), firm T cannot profitably increase f_{Tl}^* in equilibrium; otherwise firm T would increase f_{Tl}^* by a bit and let intermediaries promote product T . Hence,

$$f_{Tl}^* = p_{Tl}^* - c_T. \quad (1.2)$$

Also, $p_{Tl}^* < v_T$ does not occur in equilibrium because then firm T can profitably induce intermediary l to educate consumers and promote product T by increasing p_{Tl}^* by 2ϵ and f_{Tl}^* by ϵ for sufficiently small $\epsilon > 0$. Thus,

$$p_{Tl}^* = v_T. \quad (1.3)$$

Combining equality (1.3) and $v_D - p_{Di}^* \leq v_T - p_{Tl}^*$ for all i yields $p_{Di}^* \geq v_D$ for all i . Since consumers buy product D , their perceived utility of buying it must be non-negative. Hence,

$$p_{Di}^* = v_D. \quad (1.4)$$

Equalities (1.1) to (1.4) uniquely pin down the contracts to intermediary l : $p_{Tl}^* = v_T$, $f_{Tl}^* = v_T - c_T$, $p_{Dl}^* = v_D$, $f_{Dl}^* = N(v_T - c_T)$. Since p_{Tl}^* is the lowest product price of firm T , $p_{Di}^* = v_D$ and $f_{Di}^* = N(v_T - c_T)$ hold for all i ; otherwise firm T can profitably deviate by letting i educate consumers and promote product T . Also, $p_{Ti}^* = v_T$ and $f_{Ti}^* = v_T - c_T$ for all i ; otherwise firm D can profitably decrease its commission to i . Hence, if all intermediaries promote the deceptive product without educating consumers, then the equilibrium is unique among all deceptive equilibria: $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, $p_{Di}^* = v_D$, $f_{Di}^* = N(v_T - c_T)$ for all i .

Notice that neither firm T nor intermediaries have incentives to deviate. Firm D has an incentive to follow the above strategy when the following two conditions hold. First, firm D earns non-negative profits given the above strategies: $p_{Di}^* + \bar{a} - c_D - f_{Di}^* \geq 0$. Second, the difference of commissions is not larger than the profits from deception: $\bar{a} \geq f_{Di}^* - f_{Ti}^*$; otherwise, firm D would set $p'_{Di} = v_D - \bar{a} - \epsilon$, $f'_{Di} = f_{Ti}^* + \epsilon$ for sufficiently small $\epsilon > 0$ and let intermediaries educate consumers and promote product D . By combining these two inequalities, I obtain the following ‘‘Condition for Deception’’:

$$\min\{v_D - c_D, v_T - c_T\} + \bar{a} \geq N(v_T - c_T). \quad (\text{CD})$$

²⁵Suppose not. Then, consumers do not buy product T from intermediary l even when they are educated, and hence firm D would set $p_{Di} = v_D$ and $f_{Di} = 0$ for all i . But then firm T can profitably deviate by offering $p'_{Tl} = v_T - \epsilon$, $f'_{Tl} = \epsilon$ for small $\epsilon > 0$.

The deceptive equilibrium exists if and only if Condition (CD) holds. Notice that in this equilibrium, naive consumers' ex-post utility is $-\bar{a} < 0$, firm D earns positive profits if Condition (CD) holds with strict inequality, firm T has zero market share, and each intermediary has $1/J$ of the market share and earns $N(v_T - c_T)/J > 0$ of total profits.

In the Appendix, I show that in any deceptive equilibrium, all intermediaries promote the deceptive product and all consumers are uneducated. Also, if all consumers are educated about \bar{a} , then commissions are competed away as in Proposition 1.2. Since deceptive equilibria and non-deceptive equilibria are jointly exhaustive by definition, these considerations lead to complete characterization of the equilibria in the model:

Proposition 1.3 (Equilibria in the Model). Suppose firms market through intermediaries and all consumers are naive.

(i) A deceptive equilibrium exists if and only if Condition (CD) holds. If the deceptive equilibrium exists, then it is unique among deceptive equilibria: $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, $p_{Di}^* = v_D$, $f_{Di}^* = N(v_T - c_T)$ for all i . In the equilibrium, all consumers receive ex-post negative utility. Each intermediary promotes the deceptive product without educating consumers and earns positive profits. The deceptive firm earns positive profits if Condition (CD) holds with strict inequality. The non-deceptive firm has zero market share. Social welfare is not maximized when $v_D - c_D < v_T - c_T$.

(ii) A non-deceptive equilibrium always exists and its outcome is unique among non-deceptive equilibria. In the equilibrium, all consumers are educated, intermediaries earn zero profits, and social welfare is maximized.

By Condition (CD), holding other parameters constant, deception occurs if and only if the amount of the hidden attribute is large. In the deceptive equilibrium, each intermediary faces a key trade-off between market share and the level of commissions. On the one hand, an intermediary can increase its market share by educating consumers and attracting them from other intermediaries. On the other hand, an intermediary can earn higher commissions per customer by not educating consumers and selling the deceptive product. As a result, deception occurs if the profits from the hidden attribute allow the deceptive firm to give each intermediary a sufficiently high commission with which the transparent firm cannot compete.

If deception occurs, then having competition among intermediaries does not lower the level of commissions. This is because the deceptive firm needs to give each intermediary a high commission to maintain deception.²⁶ This result brings a new insight to the relation between commissions and the role of intermediaries: although high commissions in classical models often imply that intermediaries provide valuable or high-cost services to their

²⁶The intuition of why high commissions can be sustained under competition among intermediaries is close to Besley and Prat (2006) and Asker and Bar-Isaac (2013). Besley and Prat (2006) show that a government has an incentive to give medias sufficiently high bribes in order to prevent these medias from broadcasting bad news. Asker and Bar-Isaac (2013) show that in a repeated-game framework, a monopolistic up-stream firm may give retailers sufficiently high transfers so that no retailer would accommodate potential up-stream entrants. In these papers, however, all parties are rational and hence welfare and policy implications are different from my paper. Also, their results would be different when there are many firms or heterogenous consumers, whereas I analyze these extensions and show the robustness of my results in Section 1.6.

customers, disproportionately high commissions may indicate that intermediaries promote products in a socially-inefficient way. This result can help explain why actively-managed mutual funds and option adjustable-rate mortgages are able to profitably charge higher total prices than alternative products, such as index funds and fixed-rate mortgages.

Deception may severely harm consumer and social welfare. If Condition (CD) holds, then the deceptive firm can profitably sell an inferior product (i.e., $v_D - c_D < v_T - c_T$), leading to suboptimal social and consumer welfare. Moreover, the deceptive firm can profitably sell its product even when the product is socially wasteful (i.e., $v_D - c_D < 0$). Deception enables the survival of products that should not exist in the market.

Deception becomes less likely to occur as consumers' search intensity, N , increases. Conditional on deception, however, increasing the search intensity further raises the level of commissions. Figure 1 describes comparative statics on N when $v_D - c_D = v_T - c_T = 1$. As N increases, educating consumers becomes more attractive to each intermediary. To maintain deception, therefore, the deceptive firm must give a higher commission at the expense of own profits. Once the commission becomes so high that the deceptive firm cannot profitably maintain deception, deception is eliminated and commissions are competed down. As a result, N has a non-monotonic effect on the level of commissions.²⁷ Similarly, so long as Condition (CD) holds, the level of commissions is *increasing* in the social surplus of the transparent product ($v_T - c_T$). As an alternative product becomes more attractive, a deceptive firm needs to give higher commissions in order to maintain deception.




Figure 1.1: The equilibrium commission (f_D^*) and profits (π_D^*) as a function of consumers' search intensity (N) when $v_D - c_D = v_T - c_T = 1$.

If Condition (CD) does not hold, then all consumers are educated about the hidden attribute, intermediaries earn zero profits, and social welfare is maximized. Hence, deception is a concern when and only when consumer misperception is substantial. On the one hand, Condition (CD) holds only when $\bar{a} \geq v_T - c_T$. This indicates the lack of “minor” deception: intermediaries educate consumers about small misperceptions under competition. On the

²⁷Notice that N does not depend on the total number of intermediaries, J , but on how many intermediaries consumers visit. Section 1.5.3 discusses policies that enhance the access to intermediaries.

other hand, Condition (CD) implies that the more important the educational role of intermediaries is (the higher \bar{a} is), the less likely the intermediaries serve their role (educating consumers about \bar{a}). Intermediaries' educational role works perversely when deception is a concern. Section 1.4.3 further investigates the welfare effect on the educational role of intermediaries.

One can also think that some intermediaries care about honesty or reputation. As a natural extension of the model, consider each intermediary incurs cost $\rho > 0$ when it does not educate consumers, where ρ represents a dishonesty or reputation cost.²⁸ In this case, a deceptive firm needs to give each intermediary a commission $N(v_T - c_T) + \rho$ to maintain deception. Although the honesty concern makes deception less likely to occur, conditional on deception it further *raises* the level of commissions. Intuitively, if intermediaries incur disutility from deception, then deceptive firms need to compensate the intermediaries more. Further, I show in Section 1.6.1 that when the honesty concern fails to generate market transparency, the honesty concern can actually decrease consumer welfare through the increase of commissions.

1.4.3 Welfare Effects of Intermediaries

This subsection highlights two significant welfare effects of intermediaries under deception. Notice that if Condition (CD) does not hold, then intermediaries educate all consumers. In this case, consumers are not exploited and commissions are competed down. When Condition (CD) holds, however, perverse welfare effects arise due to the presence of expert intermediaries.

I first examine the effect on the educational role of intermediaries. To investigate it, consider an alternative case where consumers are naive and intermediaries *cannot* educate consumers about the hidden attribute. When Condition (CD) is satisfied, all consumers buy the deceptive product. Since no one can educate consumers in such a case, of course deception occurs. In this case, however, the deceptive firm can profitably decrease its commissions without the threat of education. Hence, commissions are competed down to zero in any equilibrium. The non-deceptive firm sets marginal-cost pricing as in classical Bertrand-competition models. Consumers are indifferent between the deceptive product and non-deceptive product, without taking the hidden cost into account. Hence, consumers' ex-post utility in this case is the social surplus of the alternative non-deceptive product minus the hidden cost: $(v_T - c_T) - \bar{a} < 0$. Notice that it is larger than $-\bar{a}$, i.e., the ex-post utility under deception in the model where intermediaries can educate consumers is *lower* than in the alternative case where intermediaries cannot educate.

Proposition 1.4 (Welfare Effect on the Educational Role of Intermediaries). Suppose Condition (CD) holds. Then, consumer welfare is lower when intermediaries can educate consumers about the hidden attribute than when they cannot.

²⁸See, for examples, Bolton et al. (2012) and Inderst and Ottaviani (2009, 2012c) for studies incorporating such a reputation cost.

Proposition 1.4 demonstrates that the existence of intermediaries, who can educate consumers, may decrease consumer welfare. This result indicates a perverse welfare effect on the educational role of intermediaries. To see the intuition, notice again that commissions are competed down to zero when intermediaries cannot educate consumers. To maintain deception, in contrast, high commissions are paid to intermediaries when they can educate consumers. Consumers buy the deceptive product in both cases, but in the latter case the consumers ultimately bear the cost of high commissions through the increase of the total prices of products. Hence, conditional on deception, consumer welfare is lower in the model where intermediaries can educate consumers compared to the alternative case where intermediaries cannot educate. When deception is an issue, experts may make consumers worse off due to the misalignment of educational incentives.

I next discuss the effect on the presence of competition among intermediaries. Suppose a modified model in which each consumer visits only one intermediary ($N = 1$). Since there is no competition among intermediaries, each intermediary promotes product D if and only if $v_D - c_D + \bar{a} \geq v_T - c_T$; the inequality is satisfied when Condition (CD) holds.²⁹ The equilibrium in this case is $p_{Di} = v_D$, $f_{Di} = v_T - c_T$, $p_{Ti} = v_T$, $f_{Ti} = v_T - c_T$, and each intermediary i promotes the deceptive product without educating consumers. Consumers' ex-post utility in this case is $-\bar{a}$, which is the *same* as in the model under multiple intermediaries.

Proposition 1.5 (Welfare Effect on the Presence of Competition among Intermediaries). Suppose Condition (CD) holds. Then, the ex-post utility of consumers is the same under a monopoly intermediary and under multiple intermediaries.

Proposition 1.5 sharply contrasts with the predictions from models of rational consumers with $v_D - c_D > 0$, where consumers get zero utility under a monopoly intermediary but get positive utility under competition among intermediaries. When consumers have misperceptions, having competition among intermediaries may not benefit consumers at all. The condition for deception, however, becomes stringent as N increases. Therefore, introducing competition among intermediaries in the model either makes the market transparent or does not increase consumer and social welfare.

1.5 Policy Analysis

This section discusses various policy interventions. Section 1.5.1 analyzes policies regulating commissions. Section 1.5.2 discusses direct regulations on the hidden attribute. Section 1.5.3 examines policies that lead consumers to reach more intermediaries. Section 1.5.4 discusses mandatory disclosure of commission structures.

²⁹If $v_D - c_D + \bar{a} < v_T - c_T$, then the monopolistic intermediary promotes the transparent product and consumers' ex-post utility is zero.

1.5.1 Regulating Commissions

This subsection discusses regulations on commissions. In the model, a simple intervention can eliminate deception. Suppose a policymaker caps the level of commissions. Under this regulation, intermediaries always educate consumers in order to increase market share:

Proposition 1.6 (Regulating Commissions). Suppose commissions are restricted to $f_{xi} < N(v_T - c_T)$ for all x, i . In any equilibrium, all consumers are educated about the hidden attribute, intermediaries earn zero commissions, and social welfare is maximized.

Proposition 1.6 shows that a direct price control on commissions in a competitive environment may increase welfare. Once the difference in commissions is restricted, intermediaries cannot get much higher commissions from deception. Hence, intermediaries choose to educate consumers about the hidden attribute in order to increase their market share. If Condition (CD) holds, then the ex-post utility of consumers increases from $-\bar{a}$ to $\min\{\max\{0, v_D - c_D\}, v_T - c_T\} \geq 0$ by the regulation. Social welfare also increases when $v_D - c_D < v_T - c_T$, because consumers always buy a product with a higher social surplus under the regulation.

As real-world examples, commissions were recently banned in the Australian and UK mutual-fund industries. The UK Financial Services Authority banned commissions “to address the potential for adviser remuneration to distort consumer outcomes” effective in January 2013.³⁰ Also in the pharmaceutical industry, doctors are not allowed to receive direct commissions from pharmaceutical companies in many countries. Proposition 1.6 shows that such policies can increase consumer and social welfare when deception is a concern.³¹

An alternative regulation—analogueous to a recent policy in the US mortgage industry—is to set a uniform commission in a market. In 2011, “to protect mortgage borrowers from unfair, abusive, or deceptive lending practices,” the Federal Reserve Board prohibited compensation to a mortgage broker based on terms or conditions of a mortgage transaction.³² If commissions are regulated to be uniform across products in the model ($f_{Di} = f_{Ti}$), then intermediary i has no incentive to conceal the hidden attribute. Note that this policy does not regulate the “level” of commissions because the intent of the policy is to disallow price discrimination by commissions.

As a potential advantage, regulating commissions requires policymakers to have less knowledge about hidden attributes than regulating the attributes directly. In order to regulate a product attribute itself, policymakers need to know which attributes are used by firms to exploit consumers (discussed in Section 1.5.2). In contrast, so long as policymakers know

³⁰Inducements Rules and the Retail Distribution Review Adviser Charging Rules, Financial Services Authority (October 1, 2012). Also, the Australian government banned commissions to “encourage financial advisers to become more client-focused” effective in July 2013: http://futureofadvice.treasury.gov.au/content/Content.aspx?doc=faq.htm#_What_are_the_1.

³¹Precisely, banning commissions does not necessarily predict educating consumers in the model. This is because given the regulation, intermediaries are indifferent between promoting deceptive products and promoting non-deceptive products. However, if intermediaries incur a reputation or honesty cost by not educating consumers, then—no matter how small the cost is—all consumers are educated under the regulation.

³²Banking and Consumer Regulatory Policy Press Release, Federal Reserve Board (August 16, 2011): <http://www.federalreserve.gov/newsevents/press/bcreg/20100816d.htm>.

deception is an issue in a market, they do not need to identify how firms exactly exploit consumer misperceptions in order to employ commission regulations.

Furthermore, regulating commissions has a positive effect on another relevant issue on deception: preventing firms from inventing new consumer-exploiting technologies. Suppose that before the price-setting stage, the deceptive firm is able to engage in “exploitative innovation” with a positive innovation cost $I_a > 0$ that increases the maximum hidden payment by $\Delta a > 0$. Assume that the innovation is appropriable (i.e., other firms cannot copy the innovation).³³ The next corollary highlights the positive role of intermediaries when commissions are regulated:

Corollary 1.1 (Exploitative Innovation). Suppose Condition (CD) holds in the model. Consider an extended model in which firm D has an opportunity to increase the amount of the hidden attribute from \bar{a} to $\bar{a} + \Delta a$ by paying an investment cost $I_a > 0$ prior to the price-setting stage.

(i) If there is no regulation, then firm D invests in the innovation if and only if $I_a \leq \Delta a$. Consumers’ ex-post utility is $-\bar{a} - \Delta a$ if the investment takes place and is $-\bar{a}$ otherwise. Social welfare is not maximized if $v_D - c_D < v_T - c_T$ or $I_a \leq \Delta a$.

(ii) If commissions are regulated to $f_{xi} < N(v_T - c_T)$ for all x and i , then firm D never invests in the innovation. Consumers’ ex-post utility is non-negative. Social welfare is maximized.

Corollary 1.1 (i) shows that welfare-harming innovations can occur in the absence of regulation. Since the increase in \bar{a} enables one-to-one transfer from naive consumers to deceptive firms, the deceptive firms have a strong incentive to invent a new consumer-exploiting technology. Such an investment is a pure waste from a social perspective. Moreover, it implies a vicious cycle of deception: once the hidden attribute is large enough, deception takes place, and the profits from deception further finances the development of deception, and so forth.

In contrast, Corollary 1.1 (ii) shows that deceptive firms do not invest in exploitative innovations because intermediaries would educate consumers about new hidden attributes under commission regulations. Hence, intermediaries can improve welfare through their educational role under commission regulations. So long as commissions do not distort the incentive of intermediaries, policymakers may want to have intermediaries *because of* problems of hidden attributes.

To the best of my knowledge, this is the first theoretical result of which policymakers can prevent firms from inventing unanticipated hidden attributes. Though there is a potentially huge welfare loss, this problem has not been investigated in the literature. Recently, innovations of hidden fees seem to be occurring in the credit-card, mortgage, and mutual-fund markets.³⁴ Corollary 1.1 shows a positive aspect of regulating commissions that discourages firms from inventing new hidden fees. However, intermediation does not seem to play a

³³Heidhues, Kőszegi and Murooka (2012b) investigate innovation incentives of deceptive firms in a retail market. By focusing on the appropriability of the innovation, Heidhues et al. (2012b) highlight perverse effects of innovation incentives when the up-front prices of the products are binding from below.

³⁴See, for examples, Bar-Gill and Bubb (2012), Bar-Gill (2009), and Anagol and Kim (2012).

central role specifically in the credit-card market, and a policymaker needs some other interventions to prevent deception in the market. Hence, this kind of policy works only when intermediation has a key role in a market.

One caveat regarding commission regulations is that, as Inderst and Ottaviani (2012c) and others point out, such regulations may create moral-hazard problems for intermediaries. For example, commission regulations may decrease intermediaries' incentives to search for better products for each of their customers. In contrast, exploitative innovations can be regarded as the moral-hazard problem for deceptive firms, which arises due to the possibility of deception.

1.5.2 Regulations on Hidden Attributes

This subsection discusses regulations that directly decrease the maximum amount of hidden attributes.³⁵ Notice that in the deceptive equilibrium, the profits of the deceptive firm and each intermediary are respectively $(v_D - c_D) - N(v_T - c_T) + \bar{a}$ and $N(v_T - c_T)$, and the ex-post utility of consumers is $-\bar{a}$. Hence, the decrease in \bar{a} benefits consumers one-to-one. Akin to Heidhues and Kőszegi (2010) and Heidhues et al. (2012a), this insight provides a counter-example to a popular argument against the Credit CARD Act and other consumer-protection regulations that the costs firms incur due to a regulation will be passed on to consumers. Further, a decrease in \bar{a} makes Condition (CD) less likely to hold. Once Condition (CD) is not satisfied, the market becomes non-deceptive, the level of commissions drops, and welfare is improved.

In contrast to the commission regulations described in the previous subsection, a policy decreasing \bar{a} is effective even when deceptive firms can give secret bribes to intermediaries. There are some potential drawbacks, however. It might be hard for a policymaker to identify which attributes are used to exploit consumers. Also, deceptive firms still have strong incentives to invent new hidden attributes that policymakers do not anticipate.

1.5.3 Enhancing Access to Intermediaries

As discussed in Section 1.4.2, an increase in N makes deception less likely to occur. On the one hand, Proposition 1.3 highlights the welfare increase when the number of intermediaries increases beyond a critical threshold. On the other hand, the increase in N does not affect consumer and social welfare so long as Condition (CD) holds.

It is worth mentioning that the policy increasing consumers' search intensity is robust to secret bribing and to the detailed knowledge of which attributes are hidden. However, the policy has at least one potential drawback: firms have strong incentives to invent new hidden attributes. Moreover, Section 1.6.1 and 1.6.2 show that in extended models, conditional on deception occurring the increase in N further harms naive consumers.

³⁵Although employing such regulations seem difficult in general, it may be possible in some specific cases. For example, the Credit Card Accountability, Responsibility, and Disclosure (Credit CARD) Act of 2009 limits late-payment penalties and other fees. This act prevents credit-card companies from charging high additional payments. See Bar-Gill and Bubb (2012) for detailed discussion.

Relatedly, regulations of disallowing exclusive dealings, as in the pharmaceutical industry, could be harmful to naive consumers because non-deceptive firms may not be able to sell their products under common agencies. In my model, allowing exclusive dealing is beneficial to consumers when (and only when) intermediaries affiliated with a non-deceptive firm reach a fraction of consumers, as discussed in Section 1.6.3.

1.5.4 Mandatory Disclosure of Commission Structure

In the model, naive consumers do not infer the existence of hidden attributes from product prices or commissions. If consumers can rationally anticipate the existence of hidden attributes from observing high commissions, then mandatory disclosure of commission structures is effective to eliminate deception. As a potential advantage, this policy does not require the detailed knowledge of the hidden attributes.

Evidence suggests that, however, people often do not rationally infer how the advice of experts is distorted from observable information.³⁶ Daniel, Hirshleifer and Teoh (2002) extensively discuss investor credulity in financial markets. Experimental evidence provided by Cain, Loewenstein and Moore (2005) suggests that people under-infer the strategic response of intermediaries. As empirical evidence, Malmendier and Shanthikumar (2007) show that small investors are inattentive to the systematic upward bias of stock recommendations of analysts. These investors also fail to utilize the information about affiliations of the analysts, even though affiliated analysts have a stronger upward bias than unaffiliated analysts.

Also, if consumers misinterpret the value of the products from observable information, then the disclosure of commission structure may not work well. For example, individual investors might naively guess that high commissions of mutual funds predict high performance, whereas Christoffersen et al. (2013) report that the high commissions actually predict future low performance. Finally, Section 1.6.2 shows that if such disclosure makes only a small fraction of naive consumers sophisticated and is not enough to eliminate deception, then the disclosure can decrease consumer and social welfare.

1.6 Extensions

This section analyzes extensions of the model. Section 1.6.1 investigates a model incorporating competition among deceptive firms. Section 1.6.2 examines models incorporating consumer heterogeneity in naivete. Section 1.6.3 discusses other extensions and modifications of the model.

³⁶Eyster and Rabin (2005) develop a model where a player does not rationally infer how other players' actions depend on their own situations. By applying this model, Eyster, Rabin and Vayanos (2013) analyze an asset-pricing market in which traders fail to take into account the informational content of prices.

1.6.1 Competition among Deceptive Firms

This subsection analyzes a modification of the model in which there are multiple deceptive firms as well as multiple non-deceptive firms in a market.³⁷ I focus on identifying conditions for deceptive equilibria in which each type of firm chooses the same strategy and consumers buy deceptive products. Suppose such a deceptive equilibrium exists. Note that in this model, the equilibrium profits of each deceptive firm must be zero ($f_{Di}^* = p_{Di}^* - c_D + \bar{a}$); otherwise, each of them can increase its commission by a bit and expand its market share discontinuously.

First, suppose $f_{Di}^* = 0$ for all i . Then, $p_{Di}^* = c_D - \bar{a}$. In this case, non-deceptive firms do not deviate if $v_D - c_D \geq v_T - c_T$; otherwise, they can profitably deviate by setting $p'_{Ti} = \min\{v_T - (v_D - c_D) + 2\epsilon, v_T\}$, $f'_{Ti} = \epsilon$ for sufficiently small $\epsilon > 0$. In the Appendix, I show that if $v_D - c_D > v_T - c_T$, then in any equilibrium consumers buy product D with $(f_{Di}^*, p_{Di}^*) = (0, c_D - \bar{a})$ and receive ex-post positive utility.

Second, suppose $f_{Di}^* > 0$ for some i . As in Section 1.4.2, a candidate of such a deceptive equilibrium can be pinned down to $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, $p_{Di}^* = c_D - \bar{a} + N(v_T - c_T)$, $f_{Di}^* = N(v_T - c_T)$. Notice that under these strategies, neither non-deceptive firms nor intermediaries have an incentive to deviate if Condition (CD) hold. Deceptive firms do not deviate if $v_D - c_D \leq v_T - c_T$; otherwise, they can profitably deviate by setting $p'_{Di} = c_D - \bar{a} + f'_{Di} - \epsilon$, $f'_{Di} = f_{Ti}^* - 2\epsilon$ for sufficiently small $\epsilon > 0$. In this deceptive equilibrium, no firm earns positive profits, each intermediary earns $N(v_T - c_T) > 0$ per sale, and consumers' ex-post utility is $(v_D - c_D) - N(v_T - c_T) < 0$. Proposition 1.7 summarizes these results:

Proposition 1.7 (Equilibria under Competition among Deceptive Firms). Suppose multiple firms exist for each type.

(i) If $v_D - c_D > v_T - c_T$, then in any equilibrium all intermediaries and firms earn zero profits. Consumers' ex-post utility is positive. Social welfare is maximized.

(ii) If $v_D - c_D \leq v_T - c_T$, then there exists a deceptive equilibrium in which all intermediaries earn positive profits when Condition (CD) holds. All firms earn zero profits. Consumers' ex-post utility is negative. Social welfare is not maximized if $v_D - c_D < v_T - c_T$.

Proposition 1.7 sharply illustrates the relation between profitable deception and selling inferior products. On the one hand, if deceptive products are superior to transparent products, then competition among deceptive firms leads them to decrease prices and commissions, and all profits from deception are passed back to the consumers. Neither firms nor intermediaries earn positive profits. Since all consumers buy deceptive products which are socially superior, social welfare is maximized. On the other hand, if deceptive products are inferior to transparent products, the same trade-off between the level of commissions and market share highlighted in Section 1.4.2 can still arise. It is worth emphasizing that high commissions can be kept in the equilibrium even when neither intermediaries nor firms have monopoly

³⁷The analysis does not essentially change when there are multiple deceptive firms and one non-deceptive firm.

power. Intuitively, the threat of educating consumers and promoting non-deceptive products prevents deceptive firms from decreasing commissions. Notice that consumers' ex-post utility in Proposition 1.7 (ii) is negative but larger than that in Proposition 1.3 (i). Competition among deceptive firms increases naive consumer's ex-post utility, though the utility is still negative under profitable deception.

Some empirical studies suggest a link between profitable deception and selling inferior products. In the mutual-fund industry, Gil-Bazo and Ruiz-Verdú (2009) report that mutual funds charging higher fees have worse before-fee risk-adjusted performance—product prices *negatively* reflect their valuations. Also, Del Guercio and Reuter (2012) find that actively-managed mutual funds which are recommended by financial advisors significantly underperform alternative options such as index funds.

As discussed in Section 1.4.2, the level of commissions is increasing in $(v_T - c_T)$ or N conditional on deception. In addition, when there is competition among deceptive firms, the consumers' ex-post utility is *decreasing* in the social surplus of transparent products so long as Condition (CD) is satisfied. This is because as the social surplus of non-deceptive products increases, each deceptive firm needs to give more commissions to intermediaries in order to maintain deception. This increases the level of commissions, and naive consumers ultimately bear the cost of such commissions through the increase of product prices. Similarly, the consumers' ex-post utility is *decreasing* in N so long as Condition (CD) is satisfied. Policies encouraging more access to intermediaries may hurt naive consumers if it is not sufficient to generate transparency in the market.

1.6.2 Heterogenous Consumers

This subsection analyzes markets with consumer heterogeneity in naivete. Suppose there is competition among each type of firm as in Section 1.6.1. Assume that a fraction $\sigma \in (0, 1)$ of consumers are informed as defined in Section 1.4.1: they know which products have the hidden attributes. The remaining fraction $1 - \sigma$ of consumers are naive.

Under heterogenous consumers, equilibrium outcomes depend on how intermediaries can market products. I call that an intermediary *offers* a product when the product is on the menu but is obfuscated, and that an intermediary *promotes* a product when the product is explicitly shown to consumers. I first discuss a model in which each intermediary can offer only one product at a time. I then analyze a model in which each intermediary can offer multiple products to all consumers at a time and can sort between naive and informed consumers by shrouding (i.e., hiding) non-deceptive products as well as the hidden attributes of deceptive products. I assume that $v_D - c_D \leq v_T - c_T$ throughout this subsection.³⁸

³⁸If instead a deceptive product is superior, there exists an equilibrium as the same equilibrium outcome with Proposition 1.7 (i) in each of the following cases, and both naive and informed consumers buy deceptive products in the equilibrium.

Single-Product Offer

Suppose each intermediary can offer only one product and no consumers can buy a product that is not offered by intermediaries. This single-product dealing can be regarded as a case in which firms cannot screen consumers. In a retail market, Gabaix and Laibson (2006) consider such a setting in which each firm can sell only one type of product, and hence no firm can screen between naive and sophisticated consumers ex-ante.

In this case, a candidate of a profitable deceptive equilibrium is $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, $p_{Di}^* = c_D - \bar{a} + \frac{N}{1-\sigma}(v_T - c_T)$, $f_{Di}^* = \frac{N}{1-\sigma}(v_T - c_T)$. Informed consumers do not buy the product because all intermediaries sell only deceptive products that yield negative ex-post utility. Such a deceptive equilibrium exists if the following modified condition holds:

$$(v_D - c_D) + \bar{a} \geq \frac{N}{1-\sigma}(v_T - c_T). \quad (1.5)$$

Notice that naive consumers' ex-post utility is $(v_D - c_D) - N(v_T - c_T)/(1 - \sigma) < 0$, which is *decreasing* in the fraction of informed consumers. This is because intermediaries can attract informed consumers as well as naive consumers by promoting non-deceptive products with education. Hence, as σ increases, deceptive firms need to give higher commissions to maintain deception. The increase in commissions raises the total product prices and makes naive consumers worse off. This effect might look close to the cross-subsidization effect in Gabaix and Laibson (2006), but here the effect arises even though informed consumers do not buy any product and hence do not get any benefit from the payments of naive consumers. The welfare effect of increasing informed consumers for naive consumers is non-monotonic, and is discontinuous at the threshold value at which Condition (1.5) holds with equality. Further, conditional on deception, consumer welfare is $(1 - \sigma)(v_D - c_D) - N(v_T - c_T)$ and social welfare is $(1 - \sigma)(v_D - c_D)$; both are increasing in σ if and only if the deceptive product is socially wasteful. Intuitively, since the total amount of commissions to maintain deception in the industry is independent of σ , only the fraction of consumers who take up deceptive products determines consumer and social welfare. This result implies that educational policies aimed at making consumers sophisticated to the hidden attributes can have a non-monotonic effect on welfare.³⁹

Multi-Product Offer

Next, I analyze the case in which each intermediary can offer multiple products at a time. This multiple-product dealing can be regarded as a menu contract or a multi-product marketing; Heidhues et al. (2012a) analyze such a setting in a retail market. In this case, intermediaries can shroud the existence of superior non-deceptive products as well as the hidden attributes of inferior deceptive products in order to screen consumers. Assume that

³⁹Kosfeld and Schüwer (2011) investigate a similar welfare effect of increasing sophisticated consumers. In their model, however, welfare losses come from the effort cost of educated consumers to avoid an add-on instead of socially wastefulness of products. See also Grubb (2012) for a perverse welfare effect of disclosing policies when firms screen consumers according to their tastes.

each intermediary can offer as many as products at a time but can promote only one type of product.⁴⁰ All non-promoted products are originally shrouded. Naive consumers remain ignorant of non-promoted products, and hence, can buy promoted products only. In contrast, informed consumers can buy any offered product whether the product is promoted or non-promoted. In what follows, I look for a deceptive equilibrium in which all intermediaries promote inferior deceptive products.

First, note that informed consumers buy superior non-deceptive products; otherwise, each intermediary can increase its profits by selling these products without promotion. Then, competition among intermediaries leads informed consumers to buying non-deceptive products at $(p_{T_i}^{**}, f_{T_i}^{**}) = (c_T, 0)$.

Second, if all intermediaries promote deceptive products without educating consumers, then regarding naive consumers each intermediary faces the same trade-off as in Section 1.6.1. Hence, the equilibrium outcomes for naive consumers become the same as those in Proposition 1.7. In this equilibrium, however, there is dual pricing for the non-deceptive product: non-deceptive firms offer two types of contracts, $(p_{T_i}^*, f_{T_i}^*) = (v_T, v_T - c_T)$ and $(p_{T_i}^{**}, f_{T_i}^{**}) = (c_T, 0)$, to intermediaries. In the deceptive equilibrium, intermediaries are indifferent between promoting deceptive products at $(p_{D_i}^*, v_{D_i}^*) = (c_D - \bar{a} + N(v_D - c_D), N(v_D - c_D))$ without educating consumers, and promoting non-deceptive products at $(p_{T_i}^*, f_{T_i}^*) = (v_T, v_T - c_T)$ with educating consumers. Informed consumers buy non-deceptive products, which are shrouded to naive consumers, at $(p_{T_i}^{**}, v_{T_i}^{**}) = (c_T, 0)$.

Proposition 1.8 (Equilibrium under Heterogenous Consumers and Screening). Suppose multiple firms exist for each type of product, a fraction of consumers are informed, $v_D - c_D \leq v_T - c_T$, and intermediaries can offer multiple products at a time. Then, there exists a deceptive equilibrium in which all intermediaries earn positive profits when Condition (CD) holds. All firms earn zero profits. Naive consumers receive ex-post negative utility from buying promoted deceptive products. Informed consumers receive ex-post positive utility from buying shrouded non-deceptive products. Social welfare is not maximized if $v_D - c_D < v_T - c_T$.

Intuitively, if naive consumers cannot buy products without the help of experts while informed consumers can find and buy any products, then the market is completely segregated. This result delivers a practical implication: sophisticated and naive consumers buy products at different markets or prices. Indeed, in the mutual-fund industry, some consumers buy index funds through intermediaries with paying more than 1 percent fees, whereas other consumers directly buy funds using the same index with around 0.1 percent fees. Bergstresser, Chalmers and Tufano (2009) find that broker-sold funds attain lower risk-adjusted returns than direct-sold funds do. Hackethal, Haliassos and Jappelli (2012) and Del Guercio and Reuter (2012) also find that consumers who buy products through financial advisors are

⁴⁰If an intermediary promotes both deceptive and transparent products at the same manner, then naive consumers would compare between them and notice the existence of hidden attributes. See Piccione and Spiegler (2012) for a detailed discussion of product comparability. In Section 1.6.3, I discuss a case where intermediaries promote both types of products without inducing the suspicion of naive consumers.

worse off than those who buy products directly because of commissions and operational costs.

Notice that the equilibrium outcomes would change if intermediaries were not able to shroud non-promoted products to naive consumers. In such a case, if an intermediary unshrouds, then all naive consumers buy the same products as informed consumers buy, i.e., non-deceptive products at $(p_{Ti}^{**}, f_{Ti}^{**}) = (c_T, 0)$. Hence, deceptive firms can undercut other firms by setting $(p_{Di} - \epsilon, f_{Di} - 2\epsilon)$ for small $\epsilon > 0$ without the threat of unshrouding and promoting non-deceptive products. As a result, competition drives away the profits of intermediaries and leads to $(p_{Di}^*, f_{Di}^*) = (c_D - \bar{a}, 0)$ in equilibrium. Notice that, however, the hidden attributes of deceptive products are shrouded and naive consumers still buy inferior deceptive products. This result could help explain, for example, why online search-engine companies such as Orbitz and Expedia sometimes put additional surcharges at non-salient places, although they do not seem to get high commissions from product providers. These considerations, combined with Proposition 1.8, highlight how the scope of shrouding technologies matters in markets for advice.

1.6.3 Further Extensions and Modifications

This subsection summarizes further extensions and modifications of the model. I discuss in turn a model incorporating (i) positive costs of educating consumers about hidden attributes, (ii) heterogeneity in consumers' search intensity, (iii) the possibility of vertical integration, (iv) the possibility that intermediaries can charge advising fees or give perks to consumers directly, and (v) the possibility of promoting multiple products.

Costly Education

To investigate the educational incentive of intermediaries in a clear manner, I have assumed that expert intermediaries can modify consumer misperceptions at no cost. In practice, however, educating consumers can be costly even for experts. In the Supplementary Material, I investigate an extended model in which intermediaries incur cost $\eta \geq 0$ per customer when they choose to educate. I show that if the deceptive equilibrium exists in the original model (i.e., the case $\eta = 0$), then it becomes a *unique* equilibrium in a model with any positive η . Intuitively, if some intermediary educates consumers, then other intermediaries have an incentive to free-ride because the education is costly. But then the deceptive firm would give the educating intermediary a high commission to maintain deception, and doing so is always profitable under the parameters where a deceptive equilibrium exists in the case $\eta = 0$.

In the deceptive equilibrium of such an extended model, each intermediary earns commission $N(v_T - c_T - \eta)$ per sale from the deceptive firm. Notice that as education becomes less costly (η becomes smaller), intermediaries earn higher commissions from deception. Intermediaries with more expertise earn higher commissions not because they help consumers more, but because the deceptive firm needs to give them higher commissions to maintain deception.

Heterogeneity in Consumers' Search Intensity

In the model, the number of intermediaries each consumer visits, N , is the same across consumers. Here I consider a model incorporating heterogeneity in consumers' search intensity. Let (t_1, \dots, t_J) denote the type space of consumers with associated probability distribution (q_1, \dots, q_J) . Suppose consumers with type t_s visit s number of intermediaries randomly. Then, each intermediary has measure $(s/J)q_s$ of type- t_s consumers.

Let $\tilde{N} = \sum_{s=1}^J sq_s$. If $q_1 = 0$ and Condition (CD) holds with $N = \tilde{N}$, then there exists a deceptive equilibrium in which intermediaries earn positive profits. This equilibrium outcomes are the same as in Proposition 1.3. Intuitively, so long as no intermediaries have monopoly power ($q_1 = 0$), then only the expected increase of market share from educating consumers matters in the deceptive equilibrium. If $q_1 > 0$, however, commissions in the non-deceptive equilibrium are also positive because intermediaries have monopoly power in such a case.

Vertical Integration

So far, I have assumed that firms and intermediaries are not vertically integrated. Indeed, all results are robust to allowing various kinds of such possibilities. First, note that if Condition (CD) holds, then the non-deceptive firm cannot vertically integrate with an intermediary profitably. This is because the firm has to pay more than its social surplus to buy out an intermediary. Second, if the non-deceptive firm and some intermediary are vertically integrated or form an exclusive-dealing contract a priori, then the deceptive firm has a strong incentive to buy out such an integrated intermediary. Third, if Condition (CD) holds, then the deceptive firm has an incentive to commit to disallow intermediaries from buying out products and setting prices by themselves. This is because without such a commitment, the market becomes essentially equivalent to retail markets analyzed in Section 1.4.1, and all profits from deception are competed away. Hence, the deceptive firm does not want intermediaries to set their own product prices. As examples, financial advisors and mortgage brokers are typically not allowed to change product prices (e.g., management fees and interest rates) by themselves.⁴¹

Going slightly beyond the model, a caveat for deception through intermediation is that a non-deceptive firm has an incentive to train own in-house intermediaries to educate consumers. This practice is essentially equivalent to direct marketing with education. If the cost of developing such in-house intermediaries is small, then consumers would be educated. In some industries, however, this kind of practice is either quite costly or prohibited. For example, doctors cannot sign prescription agreements with any company.

⁴¹In contrast, front-load commissions are sometimes discounted by financial advisors. See the next extension where intermediaries can charge advising fees or give perks to customers directly.

Competition on Advising Fees or Perks

So far, I have assumed that intermediaries cannot charge advising fees or give additional rebates to consumers directly. On the one hand, as described in Inderst and Ottaviani (2012a, 2012b), direct payments for advice are not prevalent in financial services. Moreover, policy regulations sometimes prevent intermediaries from charging direct advising fees or giving perks to their customers. For example, many US states prohibit life-insurance agents to charge broker fees.⁴² On the other hand, intermediaries seem to be able to set direct advising fees in some other industries.

Here I discuss how equilibrium outcomes change if intermediaries can charge direct fees or pass their profits to consumers directly. Consider a case in which intermediaries can charge and announce their advising fees to consumers after they choose which product to promote but before consumers visit them. Consumers observe these advising fees (without knowing about product attributes nor prices) and choose N intermediaries to visit simultaneously.

Suppose first that intermediaries can set only non-negative advising fees. That is, advisors can charge fees for advice but cannot give additional rewards or perks to their customers. In this case, intermediaries compete down their advising fees to zero in order to attract profitable naive consumers. Hence, none of this paper's results changes.⁴³

Suppose next that intermediaries can hand out their profits to consumers by setting negative advising fees (i.e., giving perks) upon purchase. In this case, intermediaries pass their profits to consumers through their perks. Although no intermediaries earn positive profits in equilibrium, deception through high commissions still occurs. Intuitively, intermediaries are able to give larger perks by promoting deceptive products than by promoting non-deceptive products because the intermediaries can receive higher commissions financed by deception, and naive consumers only visit intermediaries who give the largest perks. While the profits from deception are handed out to naive consumers, the deceptive firm still pays high commissions to intermediaries and naive consumers may make suboptimal purchase decisions.

As a related issue, in the US mutual-fund industry there are fee-only advisors who do not receive any commission but charge only direct advising fees to customers. Consider a modified model in which a fraction of intermediaries accept no commissions; the remaining fraction of intermediaries receive commissions and maximize profits. For simplicity, assume that such no-commission intermediaries always educate consumers and promote non-deceptive products.⁴⁴ Suppose that a deceptive firm cannot buy out such no-commission intermediaries;

⁴²See, for example, California Department of Insurance Bulletin No. 80-6.

⁴³Inderst and Ottaviani (2012c), Armstrong and Vickers (2012), and Grubb (2012) also investigate models with zero price floor. In practice, intermediaries may not be able to set negative advising fees profitably if the negative fees attract not only customers but also "arbitrageurs." The arbitrageurs are only interested in perks and can avoid any additional fees or hidden bad qualities because they do not use the product itself at all. Heidhues, Kőszegi and Murooka (2012) provide a formal model of arbitrageurs who would take advantage of overly low prices as a micro-foundation for the price floor. This model comes from Ellison's (2005) insight that firms may not want to set very-low prices because these prices would disproportionately attract less profitable consumers. Furthermore, if intermediaries give perks to consumers, then some of the consumers might become suspicious—they would think there is a catch—and try to understand how firms and intermediaries can make profits from overly low prices.

⁴⁴Alternatively, a non-deceptive firm who sells a product through own sales agency with exclusive-dealing

otherwise, the deceptive firm would vertically integrate. Also, suppose that intermediaries cannot set direct advising fees and consumers visit N number of intermediaries randomly; if naive consumers choose which intermediaries to visit based on the level of direct advising fees, then naive consumers have no incentive to visit no-commission intermediaries as discussed in the previous paragraph. In such a model, if a fraction of consumers who reach some no-commission intermediary are small, then profit-maximizing intermediaries still choose to not educate consumers about the hidden attribute. Intuitively, when most consumers are uneducated by no-commission intermediaries, the profit-maximizing intermediaries just earn profits from the remaining uneducated consumers.⁴⁵ On the other hand, if a sufficient number of consumers reach some no-commission intermediary, then profit-maximizing intermediaries also choose to educate.

Multi-Product Promotion

In the model, I assume that each intermediary can promote only one product. On the one hand, this assumption seems plausible because if an intermediary openly promotes both deceptive and transparent products to everyone, then naive consumers might compare between them and notice the existence of hidden attributes. On the other hand, in some situations intermediaries might be able to promote different types of products without inducing any suspicion of consumers. Here I discuss a modified model where each intermediary can promote both deceptive and transparent products and naive consumers never realize the existence of hidden attributes unless they are educated by some intermediary.

In this case, whether and how educated consumers buy non-deceptive products from deceiving intermediaries becomes a key. If a consumer educated by an intermediary would buy from some other deceiving intermediary when the consumer is indifferent, then the educating intermediary cannot expand its market share by educating consumers about the hidden attribute. Hence, the threat of educating consumers and promoting a transparent product is limited, and the deceptive firm can decrease its commission level without inducing the deviations of intermediaries. In practice, however, educated consumers do not seem to buy products from an intermediary who does not educate them because such a deceiving intermediary loses reputation. To capture this concept in a simple manner, assume that educated consumers discount the value of deceiving intermediaries' products by $d \geq 0$. Then, the equilibrium outcomes in Proposition 1.3 do not change for any $d > 0$. Intuitively, the trade-off between the level of commissions and market share does not change so long as educated consumers buy non-deceptive products from some educating intermediary.

Taking $d > 0$ to zero leads to a particular tie-breaking rule: if an educated consumer is indifferent between buying from an educating intermediary and a deceiving intermediary, then the educated consumer buys from the intermediary who educates her. Given this tie-breaking rule, the analysis in the model does not change by allowing the possibility of

contracts can also be regarded as such an intermediary in the model.

⁴⁵Indeed, in the US mutual-fund industry, the market share of the fee-only advisors is small. There are about 200,000 personal financial advisors in total, whereas members of fee-only personal financial advisors (NAPFA) are about 2,500.

multi-product promotion.⁴⁶

1.7 Concluding Remarks

This paper analyzes the educational incentives of intermediaries when consumers misperceive product attributes. I show that when firms can give sufficiently high commissions financed by the misperceived attributes, expert intermediaries do not educate consumers even when they are competing for their consumers. Furthermore, if it happens, then having expert intermediaries hurts consumers through the increase in commissions. When there is an appropriate regulation and commissions do not distort the incentives of intermediaries, however, expert intermediaries can work for the consumers.

In what follows, I illustrate several questions raised by, but beyond the scope of, this paper. First, except for an extension in Section 1.6.3, I focus on the case where expert intermediaries can costlessly modify consumers' misperceptions. While this assumption is useful to analyze the educational role of intermediaries in a clear manner, costs of education can be non-negligible even when consumers directly consult with experts. Indeed, recent studies report that just providing unbiased information is sometimes not enough to modify consumer misperceptions.⁴⁷ Studies by Anagol et al. (2012) and Duarte and Hastings (2012), however, show that consumers sometime (over-)react to provided information. How firms or policymakers can effectively educate naive consumers is an important topic for future research.

Second, consumers may learn about product attributes after incurring hidden costs. They may also learn from neighbors about hidden attributes. It seems that, however, merely having the opportunity of repeated sales may not be enough to eliminate deception in an emerging market.⁴⁸ Moreover, if a sufficient number of new consumers enter the market in each period,

⁴⁶A caveat is that the tie-breaking rule is not enough to sustain the equilibrium in Proposition 1.8. This is because deceptive firms would have an incentive to let uneducating intermediaries promote both (i) deceptive products with a slightly lower price and commission and (ii) non-deceptive products sold at marginal-cost pricing. This is a variant of the case where intermediaries are not able to shroud the existence of non-deceptive products, as mentioned at the last paragraph in Section 1.6.2. When $d > 0$ is sufficiently large, however, all results in Proposition 1.8 remain the same.

⁴⁷Beshears, Choi, Laibson and Madrian (2011) conduct a lab experiment on fund purchase and report that a non-negligible fraction of consumers do not take up the lowest-cost fund even when they receive all relevant information. Choi, Laibson and Madrian (2011) conduct a field experiment in which employees randomly receive either an informational survey explaining about suboptimal choices in their retirement plans or a non-informational survey. The authors find that the change of the employees' contribution rates in their retirement plans through completing the informational survey is statistically insignificant. Bhattacharya, Hackethal, Kaesler, Loos and Meyer (2012) report that mere availability of unbiased advice is not sufficient for most consumers to make the best decision.

⁴⁸To see it, suppose the market described in Section 1.3 is repeated twice and no party enters in the second period. Assume that after the first period, all consumers become informed due to some exogenous learning force. In this case, competition among intermediaries drives down commissions to zero in the second period. Given this, the trade-off between the level of commissions and market share in the first period does not change, and the deceptive equilibrium exists in the first period if Condition (CD) holds.

then deception would be sustained in every period to exploit these new consumers. A general analysis of learning and market dynamics under consumer naivete is an interesting topic.

Finally, consumers' search intensity, N , is exogenously given in this paper. All results would remain the same if consumers' visiting costs are zero for first N intermediaries and are positive after visiting N intermediaries.⁴⁹ If instead consumers incur a positive cost per visit, then all consumers would visit only one intermediary as shown in Diamond (1971). Developing a tractable consumer-search model under naivete, as well as investigating why and how naive consumers search for advice in financial markets, is left for future research.

1.8 Proofs and Supplementary Materials

1.8.1 Proofs

Proof of Proposition 1.1. In the main text. □

Proof of Proposition 1.2.

Suppose first $v_D - c_D \leq 0$. In this case, informed consumers never buy product D . In the equilibrium, $p_{Ti}^* = v_T$, $f_{Ti}^* = 0$ for all i , and all consumers buy product T .

Suppose next $v_D - c_D > 0$. Consider the case $v_D - c_D \geq v_T - c_T$; the case $0 < v_D - c_D < v_T - c_T$ can be shown by the same logic. I first show that firm D 's equilibrium profits are equal to the difference of social surplus of the products: $p_{Di}^* + \bar{a} - f_{Di}^* - c_D = (v_D - c_D) - (v_T - c_T)$. Note that firm T never sets its price below its total cost by the assumption: $p_{Ti}^* \geq c_T + f_{Ti}^*$. If $v_D - p_{Di} - \bar{a} > v_T - p_{Ti}$ and $f_{Di} > f_{Ti}$, then intermediary i never sells product T . Hence, if firm D earns profits less than $(v_D - c_D) - (v_T - c_T)$ or if an intermediary i sells product T and has positive market share, then firm D can make intermediary i promote product D by setting $p'_{Di} = p_{Di}^* - \bar{a} + v_D - v_T - \epsilon$, $f'_{Di} = f_{Di}^* + \epsilon$ for sufficiently small $\epsilon > 0$, and this ensures firm D 's profits $(v_D - c_D) - (v_T - c_T) - 2\epsilon$. By the same logic, the equilibrium profits of firm D is at most $(v_D - c_D) - (v_T - c_T)$; otherwise firm T would make intermediaries promote product T by setting $p'_{Ti} = v_T - v_D + p_{Di}^* + \bar{a} - \epsilon$, $f'_{Ti} = f_{Di}^* + \epsilon$ for sufficiently small $\epsilon > 0$.

Now I show that $f_{Di}^* = 0$ and hence $p_{Di}^* = v_D - (v_T - c_T) - \bar{a}$ for any intermediary i with positive market share. Suppose otherwise. Then, by the previous paragraph, $f_{Di}^* = p_{Di}^* - \{v_D - (v_T - c_T) - \bar{a}\} > 0$ for any i with positive market share. First, suppose some intermediary j earns zero profits. Then, firm T can profitably deviate by proposing a contract $p'_{Tj} = c_T + (1 - \epsilon)f_{Dj}^*$, $f'_{Tj} = (1 - 2\epsilon)f_{Dj}^*$ to such j with sufficiently small $\epsilon > 0$. Next, suppose all intermediaries earn positive profits by promoting product D . Then, firm D must propose its highest product price to multiple intermediaries; otherwise, firm D can profitably increase its second-highest price offered to intermediaries without inducing intermediaries' deviations. Let h be one of such intermediaries. Take firm T 's alternative contract to intermediary h such that $p'_{Th} = c_T + (1 - \epsilon)f_{Dh}^*$, $f'_{Th} = (1 - 2\epsilon)f_{Dh}^*$. For sufficiently small $\epsilon > 0$, intermediary h would promote product T —a contradiction. □

⁴⁹In a sequential consumer-search model, it is often assumed that a fraction of consumers can visit multiple shops at no cost. See, for example, Stahl (1989).

Proof of Proposition 1.3.

Note that if all consumers are educated about \bar{a} , then no intermediary earns positive profits and social welfare is maximized as in Proposition 1.2. Also, the condition for the existence of the equilibria in which all intermediaries promote product D without educating consumers is shown in the main text.

In what follows, I prove that if some consumers are uneducated about \bar{a} , then all intermediaries promote product D and all consumers are uneducated about \bar{a} . This leads to the uniqueness of the outcome among deceptive equilibria.⁵⁰ Suppose otherwise. Note that in this case, at least N number of intermediaries do not educate consumers about \bar{a} on the equilibrium path. The proof has eight steps.

(i): *Each intermediary is indifferent between promoting product D and promoting product T .* Suppose some intermediary strictly prefers to promote some product. Note that the intermediary must earn positive profits. Then, a firm providing that product can profitably decrease a commission to the intermediary by a bit—a contradiction.

(ii): *Some intermediary earns positive profits.* Suppose all intermediaries earn zero profits. Since at least N intermediaries do not educate consumers about \bar{a} , some non-educating intermediary has positive market share.

First, consider a case in which some consumers buy product D . Then, there exists intermediary i who promotes product D and has positive market share. If $p_{Di}^* > v_D - (v_T - c_T) - \bar{a}$ for some of such i , then firm T can offer $p'_{Ti} = p_{Di}^* + v_T - v_D + \bar{a}$, $f'_{Ti} = \epsilon$ for sufficiently small $\epsilon > 0$, let i educate consumers and promote product T , and increase its profits. Hence, all consumers buying product D on the equilibrium path pay $p_{Di}^* \leq v_D - (v_T - c_T) - \bar{a}$. Then, firm D can propose $p'_{Di} = v_D - (v_T - c_T) - \bar{a} + 2\epsilon$, $f'_{Di} = \epsilon$ to every intermediary who either promotes product D or does not educate consumers (or both) on the equilibrium path. Because at least N intermediaries do not educate consumers on the equilibrium path, all such intermediaries can earn positive profits by not educating consumers and promoting product D with the new contract. Also, for sufficiently small $\epsilon > 0$, no intermediary can profitably promote product T . Hence, firm D can profitably deviate—a contradiction.

Second, consider a case in which all consumers buy product T . If $p_{Ti}^* < v_T - \max\{0, v_D - c_D\}$ for some i with positive market share, then firm T can propose $p'_{Ti} = p_{Ti}^* + 2\epsilon$, $f'_{Ti} = \epsilon$ for sufficiently small $\epsilon > 0$, let i educate consumers and promote product T , and increase its profits. Hence, any intermediary i with positive market share promotes product T with $p_{Ti}^* \geq v_T - \max\{0, v_D - c_D\}$. Consider that firm D offers $p'_{Di} = v_D - \epsilon$, $f'_{Di} = \epsilon$ to all intermediaries when $v_D - c_D \leq 0$ and offers $p''_{Di} = p_{Ti}^* + v_D - v_T - \bar{a} + 2\epsilon$, $f''_{Di} = \epsilon$ to all intermediaries when $v_D - c_D > 0$. Because at least N intermediaries do not educate consumers on the equilibrium path, in either case each intermediary can earn positive profits by not educating consumers and promoting product D with the new contract. Also, for sufficiently small $\epsilon > 0$, no intermediary can profitably promote product T . Hence, firm D

⁵⁰This uniqueness result relies on an equilibrium refinement: given other intermediaries' strategies, if an intermediary chooses to educate (resp. not educate) on the equilibrium path, then the intermediary keeps choosing to educate (resp. not educate) in any other subgame whenever doing so is a best response. Without this refinement, other deceptive-equilibrium outcomes can exist due to the coordination problems of intermediaries' educational decisions.

can increase its profits—a contradiction.

(iii): *Some intermediary promotes product D.* Suppose all intermediaries promote product T . By (ii), there exists intermediary i who earns positive profits. By (i), intermediary i is indifferent between promoting product T and product D , and hence, intermediary i can earn positive profits by promoting product D with a contract (p_{Di}^*, f_{Di}^*) . It leads that $f_{Di}^* > 0$.

First, consider a case in which intermediary i would split its market share with some other intermediary if i promotes product D with a contract (p_{Di}^*, f_{Di}^*) . In this case, firm D can offer $p'_{Di} = p_{Di}^* - \epsilon$, $f'_{Di} = f_{Di}^* - 2\epsilon$ for sufficiently small $\epsilon > 0$, let i promote product D , and increase its profits—a contradiction.

Second, consider a case in which intermediary i would not split its market share with any other intermediary if i promotes product D with a contract (p_{Di}^*, f_{Di}^*) . Note that $p_{Di}^* < v_D$ must hold in this case; otherwise, only intermediary i has positive market share in the equilibrium, and firm D can profitably deviate by offering $p'_{Dj} = v_D - \epsilon$, $f'_{Dj} = f_{Di}^* - 2\epsilon$ to any intermediary $j \neq i$. But if $p_{Di}^* < v_D$, then firm D can offer $p'_{Di} = p_{Di}^* + 2\epsilon$, $f'_{Di} = f_{Di}^* + \epsilon$ for sufficiently small $\epsilon > 0$, let i promote product D , and earn positive profits—a contradiction.

(iv): *All consumers buy product D.* Suppose some consumers buy product T . Let Σ be the set of intermediaries who promote product T and have the lowest product price of product T among those who promote product T . By (i), any $s \in \Sigma$ is indifferent between promoting product D and promoting product T .

First, suppose a case in which $f_{Ts}^* = 0$ for all $s \in \Sigma$. If $p_{Ts}^* < v_T - \max\{0, v_D - c_D\}$ for some $s \in \Sigma$, then firm T can propose $p'_{Ts} = p_{Ts}^* + 2\epsilon$, $f'_{Ts} = \epsilon$ for sufficiently small $\epsilon > 0$, let s educate consumers and promote product T , and increase its profits. If $p_{Ts}^* \geq v_T - \max\{0, v_D - c_D\}$ for all $s \in \Sigma$, then firm D can offer $p'_{Ds} = p_{Ts}^* + v_D - v_T - \bar{a} + 2\epsilon$, $f'_{Ds} = \epsilon$ for small $\epsilon > 0$ to all such $s \in \Sigma$. Because each $s \in \Sigma$ can earn positive profits by promoting product D , this is a profitable deviation for firm D .

Second, suppose a case in which $f_{Ts}^* > 0$ for some $s \in \Sigma$. By (iii), some other intermediary promotes product D . Then, a fraction of consumers are indifferent between buying product T from s and buying product D from some other intermediary on the equilibrium path; otherwise, firm T can profitably deviate by offering $p'_{Ts} = p_{Ts}^* + 2\epsilon$, $f'_{Ts} = f_{Ts}^* + \epsilon$ to all $s \in \Sigma$ when $p_{Ts}^* < v_T$, and firm D can profitably deviate by offering $p'_{Ds} = v_D - \epsilon$, $f'_{Ds} = f_{Ds}^* - 2\epsilon$ to some $s \in \Sigma$ with $f_{Ts}^* > 0$ when $p_{Ts}^* = v_T$. But then, firm T can profitably deviate by offering $p'_{Ts} = p_{Ts}^* - \epsilon$, $f'_{Ts} = \max\{0, f_{Ts}^* - (N + 1)\epsilon\}$ to all $s \in \Sigma$ for sufficiently small $\epsilon > 0$ —a contradiction.

(v): *All intermediaries earn positive profits.* Suppose intermediary i earns zero profits. By (ii), let $h \neq i$ be an intermediary who earns positive profits. By (i), h is indifferent between promoting product D and promoting product T , and hence h can earn positive profits by promoting product T . Note that firm T earns zero profits by (iv). When firm T offers $p'_{Ti} = p_{Th}^*$, $f'_{Ti} = f_{Th}^*/2$ to intermediary i , then i can earn positive profits by educating consumers and promoting product T . Hence, firm T can profitably deviate—a contradiction.

(vi): *All intermediaries promote product D.* Notice that by (v), all intermediaries have positive market share on the equilibrium path. Since no consumers buy product T on the equilibrium path by (iv), all intermediaries promote product D .

(vii): *For any two contracts associated with different products, consumers are never informed after education.* Suppose otherwise: there exist contracts (p_{Di}^*, f_{Di}^*) and (p_{Tj}^*, f_{Tj}^*) such that $v_D - p_{Di}^* = v_T - p_{Tj}^* + \bar{a}$. By (i) and (v), both firms set positive commissions to all intermediaries. By (vi), intermediary i promotes product D . Consider firm T 's offer $p'_{Tj} = p_{Tj}^* - \epsilon$, $f'_{Tj} = f_{Tj}^* - 2\epsilon$. Then, intermediary j would promote product T and firm T earns positive profits—a contradiction.

(viii): *No intermediary educates consumers.* I first show that $p_{Ti}^* = v_T$ for all i . Suppose otherwise. Then by (vii), firm T can offer $p'_{Ti} = p_{Ti}^* + 2\epsilon$, $f'_{Ti} = f_{Ti}^* + \epsilon$ and let i promote product T with educating consumers—a contradiction. Furthermore, if $p_{Ti}^* = v_T$ for all i and $p_{Dj}^* < v_D$ for some j , then firm D can profitably increase its price offered to intermediary j by a bit. Hence, $p_{Di}^* = v_D$ for all i . Note that no intermediary can sell product D once consumers are educated about \bar{a} . Since each intermediary earns positive profits by promoting product D , all intermediaries choose to not educate consumers about \bar{a} . \square

Proof of Proposition 1.4.

The case in which intermediaries can educate consumers is analyzed in Proposition 1.3. When no party can educate consumers about \bar{a} , firms and intermediaries compete as in Proposition 1.2 except that consumers do not take \bar{a} into account in their purchase decisions. Hence, when Condition (CD) holds, all consumers buy product D with a price $p_{Di}^* = v_D - (v_T - c_T)$, firm T and all intermediaries earn zero profits, and firm D earns profits $v_D - c_D - (v_T - c_T) + \bar{a}$. \square

Proof of Proposition 1.5.

The case of multiple intermediaries is analyzed in Proposition 1.3. Consider a model in which all consumers visit only one intermediary. If Condition (CD) holds, then in any equilibrium the intermediary promotes product D ; otherwise, firm D can propose $p'_D = v_D - \epsilon$, $f'_D = (v_T - c_T) + \epsilon$ and let the intermediary promote product D . Also, firm D sets $p_D^* = v_D$ because otherwise firm D can profitably increase its product price by 2ϵ and its commission by ϵ . Hence, the ex-post utility of consumers under a monopoly intermediary is $-\bar{a}$. \square

Proof of Proposition 1.6.

I show that if commissions are regulated to $f_{xi} < N(v_T - c_T)$ for all x and i , then all consumers are educated in any equilibrium. I prove it by contradiction. Suppose there exists an equilibrium in which some consumers are not educated about \bar{a} . Even under the commission regulation, the proof of Proposition 1.3 still holds up to the end of (vii). Also, if some intermediary i promotes product D , then $p_{Di}^* = v_D$ and $p_{Ti}^* = v_T$ by (viii). Given this, all intermediaries promoting product D choose to not educate consumers. But because of the regulation, some intermediary would deviate to educate consumers. Hence, in any equilibrium there exist a fraction of consumers who visit one educating intermediary and $N - 1$ non-educating intermediaries. Since all intermediaries promoting product T educate consumers, each non-educating intermediary can attract consumers who only visit intermediaries promoting product D . However, firm T can offer $p'_{Ti} = v_T$, $f'_{Ti} = v_T - c_T - \epsilon$ to

non-educating intermediary i , let i promote product T , and increase its profits for sufficiently small $\epsilon > 0$ —a contradiction. \square

Proof of Corollary 1.1.

(i): Note that under the deceptive equilibrium in Proposition 1.3, all intermediaries do not educate consumers about \bar{a} and all consumers buy product D . Hence, firm D pays the innovation cost if and only if $I_a \leq \Delta a$.

(ii): Immediate from Proposition 1.6. \square

Proof of Proposition 1.7.

(i): First of all, in any equilibrium all firms earn zero profits; otherwise some firm could increase its profits by raising its commissions by a bit. Also, because $v_D - c_D > v_T - c_T$ all consumers buy product D on the equilibrium path; otherwise some type- D firm can induce intermediaries promoting product T to deviate and increase its profits.

In what follows, I show that no intermediary earns positive profits. Suppose otherwise. First, consider a case where intermediary i is indifferent between promoting product D and promoting product T with commission $f_{Ti}^* > 0$. Since $v_D - c_D > v_T - c_T$, then some type- D firm can let intermediary i deviate and make positive profits by offering $p'_{Di} = v_D - (v_T - c_T) - \bar{a} + f_{Ti}^*$, $f'_{Di} = f_{Ti}^* + \epsilon$ for sufficiently small $\epsilon > 0$ —a contradiction. Next, consider a case where intermediary i strictly prefers to promote product D . But then, a deceptive firm can profitably deviate by offering $p'_{Di} = p_{Di}^* - \epsilon$, $f'_{Di} = f_{Di}^* - 2\epsilon$ to intermediary i for sufficiently small $\epsilon > 0$ —a contradiction. Therefore, no intermediary earns positive profits if $v_D - c_D > v_T - c_T$. Since all profits—including the profits from the hidden attributes—are passed back to consumers and $v_D - c_D > v_T - c_T$, in any equilibrium consumers get ex-post positive utility and social welfare is maximized.

(ii): In the main text. \square

Proof of Proposition 1.8.

In the main text. \square

1.8.2 Further Benchmark Cases

Equilibrium When Intermediaries Are Not Necessary: Further Cases

Here I present the robustness of the benchmark result summarized in Proposition 1.1 about the specifications of timing between pricing and educating decisions. Notice that in each of the following models, there always exists a non-deceptive equilibrium in which firm T always educates consumers, a firm with lower social surplus sets its total price equal to marginal cost, and intermediaries earn zero profits.

First, consider a model in which firms first choose own prices, and after observing the prices the firms simultaneously choose whether to educate consumers. Suppose there exists an equilibrium path where consumers are not educated. Since firms are facing Bertrand-type price competition, in equilibrium each consumer is indifferent between buying product D and T without taking \bar{a} into account: $v_D - p_D^* = v_T - p_T^*$. If firm D has positive market share, then firm T can always increase its profits by educating consumers. If firm D has zero

market share, then firm T can still attract all consumers by charging a price $p'_T = p_T^* + \bar{a}/2$ in the first stage and educating consumers in the second stage. Hence, all consumers are educated in equilibrium:

Proposition 1.9 (Equilibrium When Intermediaries Are Not Necessary, Pricing-then-Education). Suppose firms directly market to consumers, choose their prices first, and then decide whether to educate consumers after observing the prices. Then, all consumers are educated in any equilibrium.

Next, consider a model in which firms first choose whether to educate consumers, and after observing the decisions the firms simultaneously set own prices. I show that if the deceptive product is socially inferior to the transparent product (i.e., $v_D - c_D < v_T - c_T$), then all consumers are educated.⁵¹ To see it, notice that $p_D^* = c_D - \bar{a}$, $p_T^* = \min\{v_T, v_T - (v_D - c_D)\}$ holds in any second-stage subgame when consumers are educated. Also, by the same logic as Proposition 1.1, either $p_D^* = c_D - \bar{a}$, $p_T^* = v_T - (v_D - c_D) - \bar{a}$ or $p_D^* = v_D - (v_T - c_T)$, $p_T^* = c_T$ holds in any second-stage subgame when consumers are uneducated. Hence, if consumers are uneducated on the equilibrium path, then in either case firm T can increase its profits by educating consumers in the first stage—a contradiction. Dahremöller (2013) shows a similar result in a market with horizontally-differentiated products.

Proposition 1.10 (Equilibrium When Intermediaries Are Not Necessary, Education-then-Pricing). Suppose firms directly market to consumers, decide whether to educate consumers first, and then choose prices after observing the educational decisions. If the deceptive product is socially inferior to the transparent product, then all consumers are educated in any equilibrium.

Equilibrium without Naivete: Further Case

In this subsection, I show that the result in Proposition 1.2 remains the same if consumers correctly anticipate the existence of a hidden attribute but do not know which product has the hidden attribute. The following result is a variant of Ellison's (2005) Proposition 4 or of Gabaix and Laibson's (2006) benchmark case where consumers are Bayesian.

Consider a case in which all consumers correctly anticipate that one firm has a hidden fee $\bar{a} > 0$, but the consumers cannot observe which firm has the hidden fee. To make a model well-defined, let $q \in (0, 1)$ be the consumers' ex-ante prior belief where firm D has the hidden fee, although the following proof and results do not depend on q . As in standard asymmetric-information models, assume that each consumer forms an ex-post belief based on Bayesian inference.

Suppose, toward to a contradiction, that there exists an equilibrium in which some consumers remain uneducated. Then, there must exist an intermediary who has positive market

⁵¹Here, a deceptive equilibrium can exist when $v_D - c_D \geq v_T - c_T$. This is because firm T cannot profitably sell its product even after education, and hence firm T has no incentive to educate consumers. In this case, however, social welfare is maximized because the deceptive product is socially superior to the transparent one.

share and does not educate consumers on the equilibrium path. Let i be such an intermediary and x be the product intermediary i promotes. Notice that uneducated consumers must form a rational belief. First, suppose a case in which uneducated consumers' belief is that intermediary i promotes product x only when product $y \neq x$ has the hidden fee. Then, firm x would let i promote product x and not educate consumers when (in terms of consumers' ex-ante perspective) product x has the hidden fee—a contradiction. Second, suppose a case in which uneducated consumers' belief is that intermediary i promotes product x in any case. Then, firm x would let i educate consumers when (in terms of consumers' ex-ante perspective) product x does not have the hidden fee, because the education would increase consumers' willingness to pay—a contradiction. Given this, consumers must form a belief such that intermediary i promotes product x only when product x has the hidden fee. Hence, uneducated consumers always correctly foresee which product has the hidden fee for sure. Then, standard Bertrand-competition arguments apply. The next proposition summarizes the result:

Proposition 1.11 (Equilibrium without Naivete: Anticipated Consumers). Suppose firms market through intermediaries and all consumers correctly anticipate the existence of \bar{a} but do not know which firm has \bar{a} . Then, in any equilibrium, all consumers correctly foresee which product has \bar{a} , and all intermediaries earn zero profits.

Intuitively, when consumers correctly anticipate the existence of a hidden attribute, education increases consumers' willingness to pay for a non-deceptive product, and hence the non-deceptive firm always induces intermediaries to educate consumers. Therefore, consumers correctly foresee that uneducating intermediaries always promote a deceptive product.

1.8.3 Costly Education

In the main text, I have assumed that intermediaries can educate consumers about hidden attributes at no cost. Although this setting allows me to investigate the educational incentive of intermediaries in a clear manner, in practice educating consumers can be costly even for expert intermediaries. This section investigates an extended model in which each intermediary incurs cost $\eta \geq 0$ per consumer when it chooses to educate the consumer. Each intermediary incurs no cost if it does not educate consumers. I show that if Condition (CD) holds in the original model—equivalent to a case $\eta = 0$ —then in a model with $\eta > 0$ the equilibrium becomes unique and is fully deceptive, i.e., all intermediaries promote the deceptive product and all consumers are uneducated.

To see the intuition, I first consider the case $N = J$. In this case, at most one intermediary educates consumers in any subgame because other intermediaries have an incentive to free-ride and avoid the education cost. But then the deceptive firm would give the educating intermediary a sufficiently high commission to employ deception ($p'_{Di} = v_D - v_T + p^*_{Ti}$, $f'_{Di} = N(v_T - c_T) - \epsilon$ for small $\epsilon > 0$), and it is always a profitable deviation given Condition (CD).

Now I prove the case $N > J$. It is straightforward to show that the fully deceptive equilibrium exists in which $p_{Di}^* = v_T$, $f_{Di}^* = N(v_T - c_T - \eta)$, $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, all intermediaries promote product D and do not educate consumers, and all consumers buy it. The proof in which there is no partial education—some consumers are educated while others are uneducated—is essentially the same as in the proof of Proposition 1.3.

In what follows, I show that there is no equilibrium in which all consumers are educated. Suppose otherwise. Notice that in such a non-deceptive equilibrium, all intermediaries earn zero net profits (profits after subtracting education costs); otherwise, a firm with positive market share would profitably undercut its prices and commissions.

In any of such equilibria, the number of intermediaries who educate consumers in any subgame is at most $J - (N - 1)$; if more than $J - (N - 1)$ intermediaries educate consumers, then all consumers are educated even when one educating intermediary does not educate, and hence some educating intermediary would profitably deviate. Also, since all consumers are educated, the number of educating intermediaries must be equal to $J - (N - 1)$. Notice that all educating intermediaries must earn positive gross profits (profits before subtracting education costs) because they incur education costs. Also, intermediaries promoting product D strictly prefer to not educate consumers in any subgame, and hence all educating intermediaries must promote product T . Then, each educating intermediary has smaller market share if it does not educate; otherwise, the intermediary would profitably deviate by promoting product T without education. It implies that there exist intermediaries who promote product D without education, and all of them would have positive market share if one educating intermediaries deviates to not educate. Also, these non-educating intermediaries earn zero profits when all consumers are educated (i.e., on the equilibrium path); otherwise, educated consumers are indifferent of between buying product D from the non-educating intermediaries and buying product T from educating intermediaries, and firm D would undercut its prices. Now, take an educating intermediary i who earns zero net profits. Then, firm D can offer $p'_{Di} = v_D - v_T + p_{Ti}^*$, $f'_{Di} = N(v_T - c_T) - \epsilon$ for sufficiently small $\epsilon > 0$. By doing so, firm D can always increase its profits by offering such an alternative contract because intermediary i earns positive net profits by promoting product D without education and Condition (CD) holds in the original model. Therefore, this is a profitable deviation from the non-deceptive equilibrium—a contradiction.

Proposition 1.12 (Uniqueness in a model with costly education). Fix all parameters other than $\eta \geq 0$, and suppose that a deceptive equilibrium exists when $\eta = 0$. Then, for any $\eta > 0$ there exists a unique equilibrium in which $p_{Di}^* = v_T$, $f_{Di}^* = N(v_T - c_T - \eta)$, $p_{Ti}^* = v_T$, $f_{Ti}^* = v_T - c_T$, all intermediaries promote product D and do not educate consumers, and all consumers buy it.

It is worth mentioning that intermediaries earn higher commissions from deception as consumers become easier to be educated (i.e., η becomes smaller). Intermediaries with more expertise earn higher commissions not because they help consumers more, but because the deceptive firm needs to give them higher commissions in order to maintain deception.

Chapter 2

Inferior Products and Profitable Deception (with Paul Heidhues and Botond Köszegi)

2.1 Introduction

In this paper, we investigate circumstances under which firms sell products by deceiving some consumers about the products' full price, focusing (in contrast to much of the literature) on deception that leads to positive equilibrium profits in seemingly competitive industries.¹ We identify a novel, perverse aspect of profitable deception: products that generate lower social surplus than the best alternative facilitate deception precisely because they would not survive in the market if consumers understood their full price, and therefore firms often make profits on exactly such bad products but not on good products.

We introduce our theory in Section 2.2, and argue that it captures a number of markets, including those for credit cards, mortgages with changing terms, managed mutual funds, life insurance, and printers, that have been invoked as conducive to hidden charges. In our model, firms are engaged in simultaneous-move price competition to sell a single homogenous product. Building on the seminal model of Gabaix and Laibson (2006), we assume that each firm charges a transparent up-front price as well as an additional price, and unless at least

¹Hidden 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.

one firm (costlessly) unshrouds the additional prices, naive consumers ignore these prices when making purchase decisions. To capture the notion that for some products, such as credit cards and mutual funds, firms cannot return all profits from later charges by lowering initial charges, we deviate from most existing work and posit that there is a floor on the up-front price.² We argue that a profitable deceptive equilibrium—wherein all firms shroud additional prices—is the most plausible equilibrium whenever it exists: it is then the unique equilibrium in the variant of our model in which unshrouding carries a cost (no matter how small), and all firms prefer it over an unshrouded-prices and therefore zero-profit equilibrium. Hence, although we characterize all equilibrium outcomes in our main propositions, we focus our analysis and discussion on identifying conditions under which a profitable deceptive equilibrium exists.

Section 2.3 presents our basic results. As two benchmark cases, we show that if the price floor is not binding, only a *zero-profit* deceptive equilibrium exists, and note that if consumers are sophisticated in that they observe and take into account additional prices, firms again earn zero profits. If the price floor is binding and consumers are naive, however, profitable deception may occur. If other firms shroud and the up-front price is at the floor, a firm cannot compete on the up-front price and can compete on the total price only if it unshrouds—but because consumers who learn of the additional prices may not buy the product, the firm may find the latter form of competition unattractive. If this is the case for all firms, a profitable deceptive equilibrium exists; and we establish that if there is a firm for which this is not the case, in equilibrium additional prices are unshrouded with probability one, and firms earn zero profits.

The above condition for a firm to find unshrouding unattractive has some potentially important implications for when profitable deception occurs. First, if the product is socially wasteful (its value to consumers is lower than its production cost), a firm that unshrouds cannot go on to profitably sell its product, so no firm ever wants to unshroud. Perversely, therefore, in a socially wasteful industry a profitable deceptive equilibrium always exists. But if the product is socially valuable, a firm that would make sufficiently low profits from deception can earn higher profits from unshrouding and capturing the entire market, so if there is such a firm only a non-deceptive, zero-profit equilibrium exists. Hence, because in an industry with many firms some firm earns low profits, entry into socially valuable industries makes these industries more transparent; and whenever deceptive practices survive in an industry with many firms, our model says that the industry is socially wasteful. Furthermore, our theory suggests a competition-impairing force in socially valuable industries that is likely to have many implications beyond the current paper: because firms face the threat that a

²In Heidhues et al. (2012), we provide a microfoundation for the price floor based on the presence of “arbitrageurs” who would take advantage of overly low prices. This microfoundation is an extreme variant of Ellison’s (2005) insight (developed in the context of add-on pricing) that firms may be reluctant to cut initial prices because these cuts disproportionately attract less profitable consumers. Another possible reason for the price floor is the consideration that if up-front prices dropped too low, consumer would become suspicious that there is a catch and not buy the product. We discuss these reasons for the price floor in Section 2.2.3 below. Ko and Williams (2011), Grubb (2012), and Armstrong and Vickers (2012) also analyze models with variants of our price-floor assumption.

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 deception.

In Section 2.4, we extend our model by assuming that there are both sophisticated and naive consumers in the market. While we find that in our single-product model sophisticated consumers can create an incentive to unshroud, we also show that in a multi-product market the situation can be radically different. In particular, if there is a superior and an inferior product, often sophisticated and naive consumers self-separate into buying the former and the latter product, respectively, and sophisticated consumers exert no pressure to unshroud the inferior product's additional prices. Worse, because the superior product renders the inferior product socially wasteful in relative terms, it *guarantees* that profitable deception in the market for the inferior product can be maintained. This observation has a striking implication: 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. As a result, firms have an incentive to push—for instance through commission-driven intermediaries or persuasive advertising—the inferior but not the superior product.

In Section 2.5, we turn to extensions and modifications of our framework. We show that if a firm has market power in the superior-product market, it may have an incentive to unshroud to attract naive consumers to itself, especially if—similarly to for instance Vanguard in the mutual-fund market—it sells mostly the superior product. But if the firm's market power is limited and unshrouding is costly, the extent to which the firm educates consumers is also limited. We also consider a specification in which there is an inefficiency associated with collecting additional prices. If the price floor is not binding and the product is socially valuable, only an unshrouded-prices equilibrium exists in this case, but if the price floor is binding or the product is socially wasteful, our qualitative results are unchanged. Finally, we discuss what happens when consumers mispredict the value of the product instead of or in addition to its price, or consumers are heterogeneous in their valuations.

In Section 2.6, we discuss the behavioral-economics and classical literatures most closely related to our paper. While a growing theoretical literature investigates how firms exploit naive consumers by charging hidden or unexpected fees, previous work has not identified the central role of wasteful and inferior products in maintaining deception and generating profits. Indeed, in most previous models competition returns all of the profits from hidden fees to consumers, so that these models cannot investigate market conditions that facilitate profitable deception. Among settings with rational consumers, our result that firms sell profitable inferior products to unknowing buyers may appear similar to the prediction of asymmetric-information models when the inferior good is cheaper to produce and consumers do not know its value. But under natural specifications these models do not predict the systematic sale of a product when an alternative that is on average superior and less profitable is available, and a rational switching-cost model also implies that a firm is better off selling a superior product.

We conclude in Section 2.7 with mentioning some policy implications of our findings,

and by pointing out important further questions raised by our model. Based on our results, we argue that deception is likely to be more prominent and have more adverse welfare consequences when the price floor is binding than when it is not binding. We also note that a policymaker can improve outcomes if she can lower additional prices, but emphasize that such policies may be limited in scope due to the difficulty of predicting or empirically establishing which price components consumers misunderstand, and of preventing firms from inventing other hidden fees.

2.2 Basic Model

In this section, we introduce our basic model of a market for potentially deceptive products. Section 2.2.1 presents our simple stylized reduced-form model, in which we impose an exogenous upper bound on the shroudable additional price and an exogenous lower bound on the transparent up-front price. In Section 2.2.2, we show how our model fits various economic applications, and how the upper bound on the additional price can arise endogenously in each case. In Section 2.2.3, we provide microfoundations for the floor on the up-front price. Table 2.1 at the end of this section summarizes how features of the applications match the elements of our formal model.

2.2.1 Setup

$N \geq 2$ firms compete for naive consumers who value each firm’s product at $v > 0$ and are looking to buy at most one item. Firms play a simultaneous-move game in which they set up-front prices f_n and additional prices $a_n \in [0, \bar{a}]$, and decide whether to costlessly unshroud the additional prices. A consumer who buys a product must pay both prices associated with that product—she cannot avoid the additional price. If all firms shroud, consumers make purchase decisions as if the total price of product n was f_n . If at least one firm unshrouds, all firms’ additional prices become known to all consumers, and consumers make purchase decisions based on the true total prices $f_n + a_n$. If consumers weakly prefer buying and are indifferent between a subset of firms, these firms split the market in proportion to exogenously given shares $s_n \in [0, 1)$.

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 equals c_{min} . In addition, we assume that $v + \bar{a} > c_n$ for all firms n ; a firm with $v + \bar{a} < c_n$ cannot profitably sell its product, so without loss of generality we can think of it as not participating in the market.

Deviating from much of the literature, we impose a floor on the up-front price: $f_n \geq \underline{f}$. We assume that $\underline{f} \leq v$; otherwise, consumers would never buy the product in equilibrium. We also assume that $\underline{f} \leq c_{min}$, so that firms are not prevented from setting a zero-profit total price.

Our main interest is in studying the Nash-equilibrium outcomes of the above game played between firms. While we fully characterize equilibrium outcomes in our main propositions,

in discussing our results we focus on conditions for and properties of deceptive equilibria—equilibria in which all firms shroud additional prices. Because no firm has an incentive to shroud if at least one firm unshrouds, there is always an unshrouded-prices equilibrium. When a deceptive equilibrium exists, however, it is more plausible than the unshrouded-prices equilibrium for a number of reasons. Most importantly, we show in Section 2.5.1 that whenever a deceptive equilibrium exists in our model with no unshrouding cost, it is the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is. In addition, a positive-profit deceptive equilibrium is preferred by all firms to an unshrouded-prices equilibrium. Finally, for the lowest-priced firms, the strategy they play in an unshrouded-prices equilibrium is weakly dominated by the strategy they play in a positive-profit deceptive equilibrium.

To simplify many of our statements as well as to facilitate comparison to classical outcomes, we define a *Bertrand outcome* as a situation in which (i) consumers buy if $v > c_{min}$ and do not buy if $v < c_{min}$; (ii) when buying consumers obtain the product at a total price of c_{min} ; and (iii) all firms earn zero profits. This is the outcome that would obtain in any equilibrium of a corresponding classical Bertrand price-competition model.

2.2.2 Applications

In this section, we discuss possible applications, especially mortgages and credit cards, of the above simple model. We consider mutual funds in the context of our multi-product model in Section 2.4 below.

Before we turn to the specific settings, we emphasize two general issues. First, we ignore the possibility that consumers are unaware not of some hidden prices, but of some advantageous product features—which a firm with the most such features always wants to explain. For instance, most consumers were initially unaware of the valuable new features of iPhone, so Steve Jobs had all the incentive to publicly demonstrate them. In many situations, however, it is some major disadvantage that consumers are unaware of. For instance, the fact that a printer can be used for printing is presumably obvious to any potential consumer, but the high cost of printing is often not.

Second, our assumption on unshrouding—that a single firm can educate all consumers about all additional prices at no cost—is in the context of most real settings unrealistically extreme, especially if (as in our credit-card application) incurring the additional prices depends partly on the consumer’s own behavior. This extreme case is theoretically useful for demonstrating that education often does not happen even if it is very easy. In Section 2.5, we show that our qualitative results survive if unshrouding is somewhat costly or reaches only a fraction of consumers. Furthermore, to the extent that unshrouding is currently very costly, our results are relevant for efforts at consumer education going forward: they indicate that even if a way to educate consumers about additional prices becomes available, we cannot rely on market forces for that education to be delivered. Finally, as we show in our earlier working paper, a model in which consumers know add-on prices but underestimate their willingness to pay for the add-on, and firms cannot educate consumers about this mistake, often generates a logic similar to that in the current model.

Mortgages with Changing Terms

Our first economic application is mortgages with changing repayment terms, where additional prices result because consumers underestimate the payments they will have to make once an initial teaser period ends. Products that fit this description include pay-option mortgages, interest-only mortgages, and mortgages with teaser rates. For instance, the Option Adjustable-Rate Mortgage (ARM) allows borrowers to pay less than the interest for a period, leading to an increase in the amount owed and sharp increases in monthly payments when the mortgage resets to an amortizing payment schedule.³ Based on existing evidence, it is plausible to assume that consumers have a limited understanding of mortgages in general, and of these mortgages in particular. Cruickshank (2000, page 127-8) finds that most consumers do not understand key mortgage features, Woodward and Hall (2012) establish that borrowers underestimate broker compensation, and Gerardi, Goette and Meier (2009) document that 30 percent of borrowers with adjustable-rate mortgages believe that they have a fixed-rate mortgage. Misunderstandings are especially likely to occur for complicated exotic mortgage products with changing terms. Indeed, evidence by Gurun et al. (2013) indicates that many borrowers are confused about mortgage reset rates, and Bucks and Pence (2008) find that borrowers who are most likely to experience large increases in payment are least likely to know their terms. And some features of these mortgages, such as the Option ARM's one- or three-*month* introductory interest rate, seem to serve only the purpose of deceiving borrowers about the product's cost.

One way to formalize the above application and type of consumer naivete is the following. Consumers can borrow a given amount for a house they value at v , and must repay their mortgage over two periods. Firm n , whose total cost of funding and administering a mortgage is c_n , offers a contract (q_n, r_n) of amounts to repay in the two periods. A naive consumer underestimates how much she has to pay in late periods, originally perceiving the total cost of repayments to be $q_n + \beta r_n$ instead of $q_n + r_n$, where $0 < \beta < 1$. Before the consumer makes her second payment, she realizes its true scale, and if in either period she perceives the total remaining payments to be greater than v , she walks away from the mortgage.

Defining f_n as the consumer's expected payment $q_n + \beta r_n$ and a_n as her unexpected additional payment $(1 - \beta)r_n$, the above formulation is equivalent to our reduced-form model from both a firm's and a consumer's point of view. Furthermore, since the consumer can walk away from the mortgage, a firm faces the constraint $r_n \leq v$, or $(1 - \beta)r_n \leq (1 - \beta)v$, which corresponds to $\bar{a} = (1 - \beta)v$ in our reduced-form model.⁴

³See "Interest-Only Mortgage Payments and Payment-Option ARMs—Are They for You?," information booklet prepared for consumers by the Board of Governors of the Federal Reserve System, available at http://www.federalreserve.gov/pubs/mortgage_interestonly/mortgage_interestonly.pdf.

⁴In a closely related formulation, consumers assume that the repayment schedule is (q_n, q_n) , i.e., that their payment will not change. In this formulation, $f_n = 2q_n$ and $a_n = r_n - q_n$, and firms face the constraint $r_n \leq v$. While (due to the last constraint) this does not match our model exactly, it generates the same logic regarding the possibility and profitability of deception.

Credit Cards

Our second economic application is credit cards, where additional prices result because consumers with a naive taste for immediate gratification borrow more than they predict or prefer. A substantial body of research in behavioral economics documents that individuals often have a taste for immediate gratification (Laibson, Repetto and Tobacman 2007, Augenblick, Niederle and Sprenger 2013, for instance), that such taste is associated with credit-card borrowing (Meier and Sprenger 2010), and that individuals may be naive about their taste (Skiba and Tobacman 2008). And consistent with the specific hypothesis that consumers underestimate costly borrowing on their credit cards, Ausubel (1991) finds that consumers are 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.⁵

We formalize the above situation using a variant of O’Donoghue and Rabin (1999). Denoting utility in period τ by u_τ , we assume that in period $t \in \{0, \dots, T\}$ a consumer makes plans to maximize total utility $u_t/\beta + \sum_{\tau=t+1}^{T+1} u_\tau$, naively believing that she will follow her currently optimal plan in the future. We suppose that the short-term discount factor β is distributed according to the cumulative distribution function Q on $[0, 1]$, where $Q(\cdot)$ is continuous and $Q(\beta)/\beta$ is bounded.⁶ We posit that the consumer’s period-0 total utility is relevant for welfare evaluations.

In period 0, the consumer decides whether to get a credit card, valuing its future convenience use at v . In each period $t = 1, \dots, T$, the consumer has an opportunity to consume 1 unit immediately, which can only be financed with a credit card. Any credit-card debt must be repaid in period $T + 1$. Consuming in period t increases u_t by 1, and repaying R in period $T + 1$ lowers u_{T+1} by R . Issuer n ’s cost of managing a consumer’s account is c_n , and all issuers acquire funds at a marginal cost of zero. An issuer chooses salient charges f_n associated with its card (such as annual fees), and a gross charge $R_n \geq 1$ for providing credit in a period; i.e., if the consumer borrows 1 in period t , she must repay R_n in period $T + 1$. To capture the notion that even salient charges are not due immediately when the consumer applies for her credit card (for instance, annual fees appear on the first statement), we assume that f_n is also paid in period $T + 1$.⁷

We now analyze the above model. From the perspective of period 0, the disutility of

⁵Beyond interest payments, credit cards can generate additional prices through various fees, including late fees and over-the-limit fees, that consumers underappreciate. Formally, the additional price resulting from these fees can be modeled identically to that from management fees in the case of mutual funds in Section 2.4 below.

⁶For presentational simplicity, we use a specification in which β divides current utility instead of the (behaviorally equivalent) more conventional specification in which β multiplies future utility. In addition, since this is irrelevant for our calculations, we do not specify whether an individual consumer’s β is correlated across periods. Our setup, for example, subsumes the case in which all consumers are naive, each consumer has given taste-for-immediate-gratification parameter β that is unknown to the firm, and the taste parameter is distributed according to Q in the population.

⁷While for simplicity we assume that all payments are made in period $T + 1$, many alternative assumptions on timing yield the same results. In particular, our analysis is unchanged so long as f_n is paid in period 1 or later, and borrowing in period t must be repaid in period $t + 1$ or later.

repayment (R_n) is greater than the utility of consumption (1) for any consumption opportunity that will arise, so in period 0 the consumer expects never to borrow on her credit card. Once she owns firm n 's credit card, however, she borrows if the current utility of consuming 1 unit ($1/\beta$) exceeds the future disutility of repayment (R_n). As a result, she borrows with probability $Q(1/R_n)$ in each period. Hence, the firm earns an expected profit of $a_n = T \cdot Q(1/R_n)(R_n - 1)$ from the consumer's borrowing. From the perspective of period 0, each instance of borrowing induces a loss of $R_n - 1$ in total utility, so the consumer loses exactly as much as the firm gains. The maximum ex-post profit the firm can earn is $\bar{a} = T \max_R Q(1/R)(R - 1)$, which exists by our assumptions on $Q(\cdot)$.

Other Applications

We discuss several other applications more briefly and informally.

Printers. In this application, discussed for instance by Gabaix and Laibson (2006), additional prices are generated by consumers' failure to anticipate high cartridge prices. For instance, Hall (1997) and the UK's Office of Fair Trading report that a large majority of consumers buying a printer do not know the price of a cartridge. To formalize this situation in the context of our model, we can assume that the up-front price is the printer price, and that when buying a printer, consumers act as if cartridges will cost nothing. Once a consumer has her printer, she realizes how much cartridges cost, and has a downward-sloping demand for printing that puts a bound on a firm's ex-post profits. A high cartridge price generates a utility loss for the consumer and profits for the firm, and hence can be thought of as an additional price.⁸

Other Products with Add-Ons. Some other products with add-ons, including bank accounts and hotels, have also been invoked by researchers as examples of potentially deceptive products.⁹ In these applications, consumers buy a base product—account maintenance in the case of bank accounts and a room in the case of hotels—and can then buy an add-on—such as costly overdrafting in the case of bank accounts or gambling entertainment in the case of Las Vegas hotels—whose cost they do not appreciate at the moment of purchase. We can think of the price of the base product as the up-front price and the utility loss from unexpectedly high add-on payments as the additional price. A consumer's underestimation of additional expenditures can—in a setting formally similar to that for printers—be due to her underestimation of add-on charges (e.g., overdraft fees), or—in a setting formally similar to that for credit cards—to her underestimation of her own ex-post demand (e.g., gambling in Las Vegas). Firms can impose additional prices by charging high prices for add-ons, but these profits are limited by consumers' ex-post demand response to add-on prices.

Life Insurance. Gottlieb and Smetters (2012) document that premiums for life insurance

⁸An economically important feature of this application is that consumers' demand response to cartridge prices not only limits firms' ex-post profits, but—because some valuable trades are not taking place—also generates an inefficiency. This point is related to Shapiro's (1995) insight that high aftermarket prices can create a dead-weight loss even if competition leads to low prices in the primary market. While we ignore this inefficiency in our basic model, we identify its implications for our theory in Section 2.5 below.

⁹See Armstrong and Vickers (2012) for bank accounts and Gabaix and Laibson (2006) for hotels.

in the US are above marginal cost early in life but below marginal cost late in life, and argue that the most plausible way to understand this pattern is to assume that consumers underappreciate the possibility of lapsing when purchasing their life-insurance policy. We can interpret these findings in our framework by positing that the expected premium payment corresponds to the up-front price, and unexpected lapsing—which lowers a consumer’s utility through the loss of valuable insurance but increases the insurer’s profits due to the below-marginal-cost pricing at that stage—corresponds to the additional price. The profit from such a “lapse-based” additional price is limited, for instance, by the consideration that overly front-loaded prices would lead fewer consumers to lapse.

2.2.3 Microfoundations for the Floor on the Up-Front Price

Our model exogenously imposes a floor on the up-front price. In this section, we discuss economic forces that could endogenously generate such a constraint.

Adverse selection. The price floor captures in an extreme form the idea that firms might be reluctant to cut prices too low because doing so would disproportionately attract less profitable consumers. In Heidhues et al. (2012), we provide one such adverse-selection-based microfoundation for a price floor. In our model, there is a share of “arbitrageurs” who at a cost e can enter the market, and who avoid the additional price because they are not interested in using the product itself. We show that if the share of arbitrageurs is sufficiently large—or if arbitrageurs can buy multiple units of the product—then firms do not cut their up-front price below $-e$: if a firm did so, arbitrageurs would enter and render the price cut unprofitable to the firm. Firms therefore behave as if they were facing a price floor of $\underline{f} = -e$.

Since in many instances the cost of arbitrage appears to be small (i.e., $e \approx 0$), the arbitrageur argument often implies a price floor close to zero. In this case, our price floor is nearly identical to a no-negative-prices constraint that has been imposed by some previous researchers.¹⁰ A price floor of about zero likely applies to credit cards and bank accounts as defined in Section 2.2.2, and it is plausible to assume that this price floor is often binding. As a simple illustration in a specific case, consider the finding of Hackethal et al. (2010) that German bank revenues from security transactions amount to €2,560 per customer per year (2.43% of mean portfolio value). If a bank handed out such sums ex ante—even if it did so net of account maintenance costs—many individuals would sign up for (and then not use) bank accounts just to get the handouts.

Furthermore, as we discuss in Heidhues et al. (2012) and as formally shown by Ko and Williams (2011), in some situations an adverse-selection consideration can generate a positive floor on the up-front price. Suppose that there are sophisticated consumers who can avoid the additional price, and who can buy an alternative product that gives them a consumer surplus $s \geq 0$. If $s > v - c_{min}$ and the share of sophisticated consumers is sufficiently

¹⁰For instance, Armstrong and Vickers (2012) and Grubb (2012) impose such a condition in the context of a market with naive consumers, and Farrell and Klemperer (2007) discuss the same constraint in the context of switching-cost models.

large, then $f \geq v - s$ must hold: if a firm charged a lower up-front price, it would attract sophisticated consumers and lose money. As a result, firms act as if they were facing a price floor of $v - s$. For instance, while many consumers who visit Las Vegas end up gambling a lot in their hotel, if a hotel became exceedingly cheap relative to its quality, it might attract vacationers uninterested in and untempted by gambling, and hence not get the additional price.

Of course, the way we model the threat of arbitrage—that a firm does not attract any arbitrageurs if it sets an up-front price above a bright-line number, but faces a flood of arbitrageurs if it sets a lower up-front price—is extreme. The intuitions for our main qualitative results on the role of wasteful and inferior products in profitable deception require only that firms are less willing to cut the up-front price than a transparent total price, so that they make higher profits with shrouded than with unshrouded additional prices. Even if the behavior of arbitrageurs is not as stark as in our model, this is the case if—similarly to the model of add-on pricing by Ellison (2005)—firms cannot avoid adverse selection when cutting their up-front price, but can avoid adverse selection when cutting their transparent total price.

Suspicion. Consumers’ suspicion that “there might be a catch” can also impose a price floor. Intuitively, if consumers see prices that they consider too low for firms to break even, they might deduce that firms must be charging additional prices as well. Since firms want to prevent consumers from reaching such a conclusion, they avoid up-front prices that seem unrealistically low for the product they are offering.

Although a full formal analysis of consumer suspicion and its implications for pricing is beyond the scope of this paper, we provide a game-theoretic formalization of the above argument. In order to avoid going through the main analysis of Section 2.3 at this stage, we suppose that firms cannot unshroud, demonstrating how the floor on the up-front price arises in a deceptive equilibrium of our main model.¹¹ Consumers initially believe that firm n is restricted to charging $a_n = 0$ with probability $1 - \epsilon$, and can choose $a_n \in [0, \bar{a}]$ freely with probability ϵ . The probabilities of different firms being “deceptive” types are independent. Consumers also believe that firms’ marginal cost is $\tilde{c} \in (c - \bar{a}, v)$, and that firms are playing a perfect Bayesian equilibrium with these parameters. In reality, firms are playing Nash equilibrium in a simultaneous-move game in which they choose $f_n \in \mathbb{R}$ and $a_n \in [0, \bar{a}]$ with true costs c , and take consumers’ behavior as given.

We argue that if $\tilde{c} + \epsilon\bar{a} \leq \min\{c, v\}$, then there is an equilibrium in which all firms choose an up-front price of \tilde{c} , so that \tilde{c} acts as a floor on the up-front price. In this equilibrium, firm n charges $f_n = \tilde{c}$, $a_n = \bar{a}$, and consumers believe that firm n offers the contract $f_n = \tilde{c}$, $a_n = 0$ if restricted to $a_n = 0$ and the contract $f_n = \tilde{c}$, $a_n = \bar{a}$ if not restricted to $a_n = 0$. If firm n

¹¹Since in a deceptive equilibrium firms find unshrouding unattractive, our analysis below showing that firms act as if they were facing a price floor applies to any candidate positive-profit deceptive equilibrium even if firms can unshroud. Furthermore, note that whenever a firm unshrouds, all consumers become informed about additional prices, so both the suspicion issue and—since a firm can choose a total price below cost—the floor become inconsequential. As a result, the conditions for a firm to find unshrouding unattractive, and therefore also the conditions for a deceptive equilibrium to exist, are identical with an endogenous suspicion-based price floor and an exogenous price floor.

Table 2.1: Summary of Main Applications of Single-Product Model

Economic application	Elements of our model		
	Up-front price	Source of additional price	Source of price floor
Mortgages	Expected total payment	Post-reset payment	Suspicion?
Credit cards	Annual fee	Interest payments	Adverse selection
Bank accounts	Maintenance fee	Overdraft charges	Adverse selection
Life insurance	Expected total premium	Lapsing	None
Printers	Printer price	Cartridge price	Suspicion?
Casino hotels	Room price	Gambling	Adverse selection?

The table summarizes what the key elements of our formal model correspond to in the context of our applications. We indicate with a question mark when a price-floor-generating mechanism is plausible, but determining with confidence whether it induces a binding price floor in the given application is difficult. For the additional price, we indicate the source in each case; the actual additional price, \bar{a} , is the utility loss the consumer suffers from incorrect expectations about the given outcome.

charges $f_n \geq \tilde{c}$, consumers maintain their prior beliefs regarding whether the firm is restricted to charging $a_n = 0$. If firm n charges $f_n < \tilde{c}$, consumers—thinking that without additional prices the firm could not survive—believe with probability one that it charges $a_n = \bar{a}$.¹² Therefore, firm n cannot profitably attract consumers if the above condition holds.¹³

Note two interesting and potentially important aspects of a suspicion-generated price floor. First, profitable deception can occur exactly *because* consumers are suspicious of the possibility of deception. Second, the argument implies that if consumers have beliefs about costs that are close to realistic (i.e., \tilde{c} is close to c), then little of the profits from the additional price are competed away ex ante, implying large profits from deception.

We also emphasize, however, that the scope for suspicion to generate a price floor is difficult to assess. Since we do not know what consumers believe about firms' costs and pricing strategies—both crucial ingredients of the argument—we do not know whether a floor applies in any specific case. Nevertheless, we feel that the argument is intuitive and possibly relevant in some of our applications. In particular, while printer buyers are prone not to pay attention to cartridge prices, if printers were exceedingly cheap, they might think more about how firms make their money. And while mortgage borrowers are prone not to consider post-reset terms, if the initial payment was too low, they might look for the catch in their contract.

¹²This inference is implied, for instance, by the dominance test of Cho and Kreps (1987, page 199).

¹³As the consumers correctly foresee that firm n charges \bar{a} if it lowers the upfront price below \tilde{c} , firm n needs to charge a total price below $\tilde{c} + \epsilon\bar{a}$ to outcompete its competitors. Hence, if $\tilde{c} + \epsilon\bar{a} \leq c$ it is impossible to profitably attract consumers by undercutting competitors' up-front price.

2.3 Profitable Deception

This section analyzes our basic model. We start in Section 2.3.1 with two benchmarks in which firms cannot earn positive profits by shrouding. In Section 2.3.2, we turn to our main result: we establish conditions under which profitable deception can be maintained. In Section 2.3.3, we discuss the role of social wastefulness in facilitating deception.

2.3.1 Benchmarks: Non-Binding Price Floor or Sophisticated Consumers

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

Proposition 2.1 (Equilibrium with Non-Binding Price Floor). Suppose $\underline{f} \leq c_{min} - \bar{a}$. For any $\psi \in [0, 1]$, there exist equilibria in which unshrouding occurs with probability ψ . If shrouding occurs in an equilibrium, firms earn zero profits, and consumers buy the product from a most-efficient firm, pay $f = c_{min} - \bar{a}$, $a = \bar{a}$, and get utility $v - c_{min}$. If unshrouding occurs in an equilibrium, the Bertrand outcome obtains.

Proposition 2.1 implies that if the price floor is not binding, there is always a deceptive equilibrium. Since in a deceptive equilibrium consumers do not take into account additional prices when choosing a product, firms set the highest possible additional price, making existing consumers valuable. Similarly to the logic of Lal and Matutes's (1994) loss-leader model as well as that of many switching-cost (Farrell and Klemperer 2007) and 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, no firm has an incentive to unshroud. Note that the deceptive equilibrium may be socially inefficient: since consumers do not anticipate additional prices, they can be induced to buy a product whose value is below production cost (i.e., $v - c_{min}$ can be negative).

By Proposition 2.1, any equilibrium outcome is identical either to that in the above deceptive equilibrium, or to that in an unshrouded-prices equilibrium, in which unshrouding occurs with probability one. As we have mentioned, since unshrouding is costless, if a firm unshrouds it is (weakly) optimal for other firms to unshroud as well. If unshrouding occurs, consumers make purchase decisions based on the total price, so firms effectively play a Bertrand price-competition game, leading to a Bertrand outcome.

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

Proposition 2.2 (Equilibrium with Sophisticated Consumers). Suppose all consumers are sophisticated. Then, the Bertrand outcome obtains.

Since sophisticated consumers understand the total price, firms cannot break even by selling a socially wasteful product (for which $v < c_{min}$). Furthermore, by standard Bertrand-competition logic, firms make zero profits in selling a socially valuable product as well. For this outcome to obtain, our assumption that $\underline{f} \leq c_{min}$ is crucial.

2.3.2 Naive Consumers with a Binding Price Floor

Taken together, Propositions 2.1 and 2.2 imply that for profitable deception to occur in our model, 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 $\underline{f} > c_n - \bar{a}$ for all n . This condition means that a firm can (if it is able to set a sufficiently high additional price) make profits when setting an up-front price equal to the floor.

We first identify sufficient conditions for a deceptive equilibrium to exist. As above, in such an equilibrium all firms set the maximum additional price \bar{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 $s_n(\underline{f} + \bar{a} - c_n)$. For this to be an equilibrium, no firm should want to unshroud additional prices. If unshrouding occurs, consumers are willing to pay exactly v for the product, so firm n can make profits of at most $v - c_n$ by unshrouding. Hence, unshrouding is unprofitable for firm n if the following ‘‘Shrouding Condition’’ holds:

$$s_n(\underline{f} + \bar{a} - c_n) \geq v - c_n. \quad (\text{SC})$$

A deceptive equilibrium exists if (SC) holds for all n . Furthermore, since $s_n < 1$, Condition (SC) implies that $\underline{f} + \bar{a} > v$, so that in a deceptive equilibrium consumers receive negative utility. Proposition 2.3 summarizes this result, and also states that if some firm violates Condition (SC), there is no deception in equilibrium:¹⁴

Proposition 2.3 (Equilibrium with Binding Price Floor). Suppose $\underline{f} > c_n - \bar{a}$ for all n .

I. If Condition (SC) holds with a strict inequality for all n , shrouding occurs either with probability one or with probability zero. In the former case, all firms offer the contract $(f_n, a_n) = (\underline{f}, \bar{a})$ with probability one, consumers receive negative utility, and firms earn positive profits. In the latter case, the Bertrand outcome obtains.

II. If Condition (SC) is violated for some n , unshrouding occurs with probability one and the Bertrand outcome obtains.

Part I of Proposition 2.3 identifies conditions under which a profitable deceptive equilibrium exists. The intuition for why firms might earn positive profits despite facing Bertrand-type price competition is in two parts. First, as in previous models and as in our model with a non-binding price floor, firms make positive profits from the additional price, and

¹⁴We ignore the knife-edge case in which no firm violates Condition (SC) but the Condition (SC) holds with equality for some firm n' . The equilibrium outcomes characterized in Case I of Proposition 2.3 remain equilibrium outcomes in this case, but there can be additional equilibrium outcomes.

to obtain these ex-post profits each firm wants to compete for consumers by offering better up-front terms. But once firms hit the price floor, they exhaust this form of competition without competing away all ex-post profits.

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

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

Part I of Proposition 2.3 also states that whenever a deceptive equilibrium exists, there is also an equilibrium in which unshrouding occurs with probability one and—because prices are then transparent and the standard Bertrand logic holds—the Bertrand outcome obtains. As in the case of a non-binding price floor (Proposition 2.1), since unshrouding is costless, if a firm unshrouds it is (weakly) optimal for other firms to unshroud as well. Unlike in the case of a non-binding price floor, however, there is no equilibrium in which unshrouding occurs with an interior probability.¹⁵

Part II of Proposition 2.3 establishes that Condition (SC) is not only sufficient, but also necessary for profitable deception to occur: if the condition is violated for some firm, then the additional prices are unshrouded with probability one and the Bertrand outcome occurs. To understand the rough logic, assume toward a contradiction that all firms shroud with positive probability. Then, a firm can ensure positive profits by shrouding and choosing prices (\underline{f}, \bar{a}) , so that in equilibrium firms earn positive expected profits. Since a firm that sets the highest total price when unshrouding has zero market share if some other firm unshrouds, to earn positive profits it must be that with positive probability all rivals set higher total prices when shrouding. Whenever a firm sets one of these high total prices, the only event in which it can make profits is when shrouding occurs, so we can think of its incentives by conditioning on this event. Doing so, arguments akin to those above imply that in this high range firms set prices (\underline{f}, \bar{a}) , and hence firm n earns $s_n(\underline{f} + \bar{a} - c_n)$. But firm n can earn $v - c_n$ by unshrouding, so a firm that violates Condition (SC) prefers to unshroud.

¹⁵Our proof of this claim elaborates on the following contradiction argument. Suppose that shrouding occurs with an interior probability, and consider a firm that unshrouds with positive probability and sets the highest price of any firm conditional on unshrouding. When setting this highest price, the firm earns profits only if all other firms shroud, and even then it earns at most $v - c_n$. If it instead shrouds and sets (\underline{f}, \bar{a}) , then it earns at least $s_n(\underline{f} + \bar{a}_n - c_n)$ if all other firms also shroud. By Condition (SC), therefore, the firm prefers to shroud with probability 1, a contradiction.

2.3.3 Socially Valuable versus Socially Wasteful Products

We now use Proposition 2.3 to identify some circumstances under which profitable deception does versus does not occur. As we have argued in Section 2.2, whenever a deceptive equilibrium exists, it is the only plausible equilibrium, so our discussion assumes that in this situation it will always be played. We distinguish two cases according to whether firms produce a socially valuable or a socially wasteful product.¹⁶

Non-vanishingly socially valuable product (there is an $\epsilon > 0$ such that $v > c_n + \epsilon$ for all n). In this case, the right-hand side of Condition (SC) is positive and bounded away from zero. Hence, a firm with a sufficiently low s_n violates Condition (SC)—since it earns low profits from deception, it prefers to attract consumers through unshrouding—and thereby guarantees that only a zero-profit unshrouded-prices equilibrium exists. This implies that a deceptive equilibrium can exist if the number of firms is sufficiently small, but not if the number of firms is large—as some firm will then violate Condition (SC). Furthermore, an increase in the number of firms can induce a regime shift from a high-price, deceptive equilibrium to a low-price, transparent equilibrium.¹⁷

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

Corollary 2.1 (Wasteful Products). Suppose $\underline{f} > c_{min} - \bar{a}$ and $v < c_n$ for all n . Then, a profitable deceptive equilibrium exists.

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

The mortgages with changing repayment terms we have discussed in Section 2.2.2 might be a good example for this case of our model. While a sharply increasing payment schedule may make sense for consumers who confidently expect drastic 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 chose mortgages with such schedules. Furthermore, by leading many consumers to overborrow and get into financial trouble, these products might well have lowered social welfare. At the same time, some of them were widespread and lucrative: Option ARMs, for instance, represented 19 percent of Countrywide’s (the then-largest lender’s) originations in 2005, and 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). Our model says that such

¹⁶To avoid unenlightening additional cases, we do not discuss in-between situations in which the product is valuable to produce by some firms but not other firms.

¹⁷The reason for stating our result for *non-vanishingly* socially valuable products is that if the social value of a firm’s product ($v - c_n$) could be arbitrarily small, then even a firm with low profits from deception might not be willing to unshroud. A precise condition for when unshrouding must occur is $N > (\underline{f} + \bar{a})/\epsilon$. Then, $s_n < \epsilon/(\underline{f} + \bar{a})$ for some n , and for this n we have $s_n(\underline{f} + \bar{a} - c_n) < \epsilon < v - c_n$, in violation of Condition (SC).

mortgages continued to be sold and remained profitable in a seemingly competitive market not despite, but *exactly because* they were socially wasteful.

The different logic of socially valuable and socially wasteful industries in our model yields two potentially important further points. First, our theory implies that if an industry experiences a lot of entry and does not come clean in its practices, it is likely to be a socially wasteful industry. Second, our theory suggests a general competition-impairing feature in valuable industries that is not present in wasteful industries: to reduce the motive to deviate from their preferred positive-profit deceptive equilibrium in a valuable industry, each firm wants to make sure competitors earn sufficient profits from shrouding. This feature is likely to have many implications beyond the current paper (as, for instance, for innovation incentives in Heidhues et al. 2012), and implies that wasteful industries may sometimes be more fiercely competitive than valuable ones.

2.4 Sophisticated Consumers and Multi-Product Markets

Our analysis has so far assumed that all consumers are naive. In this section, we discuss the implications of assuming that some consumers are sophisticated in that they observe and take into account additional prices when making purchase decisions. We begin in Section 2.4.1 by pointing out how this change modifies the logic of our basic model, and then consider a multi-product market in Section 2.4.2. Throughout this section, we assume that the proportion of sophisticated consumers is $\kappa \in (0, 1)$, and that the price floor is binding: $\underline{f} > c_n - \bar{a}$ for all n .

2.4.1 Sophisticated Consumers in Our Basic Model

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

$$(1 - \kappa)s_n(\underline{f} + \bar{a} - c_n) \geq v - c_n, \quad (2.1)$$

and a deceptive equilibrium exists if and only if Condition (2.1) holds for all n .

Condition (2.1) for the existence of a deceptive equilibrium has two notable implications. First, if the product is socially wasteful, then as before a profitable deceptive equilibrium always exists. The reason is simple: sophisticated consumers never buy a socially wasteful product in equilibrium, so their presence is irrelevant—firms just exploit naive consumers. But second, if the product is socially valuable, the condition for a deceptive equilibrium to

exist is stricter if sophisticated consumers are present than if they are not. Intuitively, while sophisticated 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.

2.4.2 Sophisticated Consumers with an Alternative Transparent Product

Setup. We modify our model above by assuming that each firm has a transparent product in addition to the potentially deceptive product we have described. Consumers' value for the transparent product is $w > 0$, and firm n 's cost of producing it is $c_n^w \geq 0$. We let $\min_n \{c_n^w\} = c_{min}^w$, and assume that there are at least two firms whose cost of producing product w is c_{min}^w . Crucially, we posit that product w is socially valuable ($w - c_{min}^w > 0$), and is not inferior to product v : $w - c_{min}^w \geq v - c_{min}$. Consumers are interested in buying at most one product. Firms simultaneously set the up-front and additional prices for product v , the single transparent price for product w , and decide whether to unshroud the additional prices of product v . If consumers weakly prefer buying and are indifferent between a number of firms in the market for product v or w , firms split the respective market in proportion to s_n or s_n^w , respectively; and if consumers are indifferent between products v and w , a given positive fraction of them chooses product w . We define a Bertrand outcome as a situation in which consumers buy a social-surplus-maximizing product at a total price that equals the product's lowest marginal cost, and firms earn zero profits; this is the outcome that would obtain in classical Bertrand price competition.

Applications. Our analysis is motivated by the observation that in many markets for deceptive products, alternatives that are more transparent than and arguably superior to the deceptive products exist. For instance, while some borrowers with bad credit ratings may have been offered only exotic mortgages with changing repayment terms (product v), many borrowers also had access to simple traditional mortgages (product w), and indeed are likely to have been served better by these products. One way to think about this distinction in our reduced-form model is to assume that the traditional mortgage has higher gross value due to the superior consumption-smoothing properties of its flat repayment schedule ($w > v$), but it costs firms the same to provide. Similarly, a typical consumer who uses a credit card (product v) could use a low-cost debit card (product w) for the same set of basic services, and come out ahead by avoiding large fees and high interest payments. In this case, however, it is the inferior product that has the higher gross value because of perks such as airline miles associated with it ($v > w$), although the value these perks create for consumers are not greater than the extra cost to firms.¹⁸

¹⁸Perks such as airline miles or free rental-car insurance do not seem natural complements of credit services, so tying these together with credit cards is unlikely to generate an efficiency gain. Hence, it would be as or more efficient for consumers not to get these perks, or to obtain them elsewhere. As we discuss below, our model naturally explains why credit-card issuers might offer inefficient perks: this provides a way to compete for consumers when the price floor is binding. See especially Footnote 24.

As an additional application that we consider to be a leading example for our multi-product model, we discuss mutual funds. Although not everyone agrees with this view, many researchers believe that because few mutual-fund managers can persistently and significantly outperform the market, most managed funds are inferior to index funds.¹⁹ In addition, different types of empirical evidence suggest that the recurring management fees of managed funds are at least partly shrouded to consumers, and hence can be thought of as an additional price. First, using a natural experiment in India, Anagol and Kim (2012) show that mutual-fund investors are less sensitive to amortized “initial issue expenses” than to otherwise identical “entry loads” paid up front. Second, although such a relationship is difficult to interpret due to factors such as advertising and marketing, Barber, Odean and Zheng (2005) and Bergstresser et al. (2009) report that net fund flows are *increasing* in total annual fees.

Formally, we think of a low-cost index fund as the superior product, and of a high-cost managed fund as the inferior product. We let v and w be consumers’ respective values from the funds before expenses, and assume that $v > w$. There are two reasons that consumers’ valuation for a managed fund might be higher. First, it might be that managed funds produce higher gross returns, even though they produce lower returns net of expenses. Second, similarly to the notion of “money doctors” in Gennaioli et al. (2012), investors might derive utility from knowing that someone is looking after their money. While in this section we suppose that naive consumers correctly perceive v , we point out in Section 2.5.3 that our results on the multi-product market are essentially unchanged if—consistent with the notion that investors overestimate the return of managed funds—naive consumers overestimate v .

Let f be the salient up-front fees, such as salient front loads, associated with a mutual fund. A managed fund selects management and other hidden fees $\hat{a} \geq 0$, generating net return $v - \hat{a}$. After a holding period of $q(\hat{a})$, the consumer realizes that the fund has low returns, sells the fund, and buys an index fund. We assume that $q(\cdot)$ is continuous and $q(\hat{a})\hat{a}$ is bounded. To map this setting to our model, we let $a = q(\hat{a})\hat{a}$, and set $\bar{a} = \max_{\hat{a}} q(\hat{a})\hat{a}$, which exists by our assumptions on $q(\cdot)$. As an extreme way of capturing the fact that the expenses and fees of most index funds are very low, we assume that $c_n^w = 0$, and that index funds charge no fees. Once we allow for misperception of v , therefore, this formulation allows for two kinds of mistakes by naive consumers: overestimating the returns a fund manager can generate, and underestimating the impact of yearly management fees on fund returns.

To complete our discussion of this application, we argue that legal constraints on mutual funds create a price floor $\bar{f} = 0$. Note that in order for a firm to charge a negative up-front price, it would effectively have to pay new investors, for instance in the form of cash back, discounts, or other incentives, for buying fund shares. But 15 USC § 80a-22d of the U.S. Code (part of the chapter regulating investment companies and their advisers) prohibits mutual funds from selling shares at a price different from the public price described in the

¹⁹See, for example, Carhart (1997) and Kosowski, Timmermann, Wermers and White (2006). Berk and Green (2004) and Berk and van Binsbergen (2012), however, argue that the evidence is consistent with the hypothesis that most managers can outperform the market, and that managed funds are not inferior to index funds.

prospectus, so that price discrimination between old and new investors is impossible.^{20,21}

Analysis. We now turn to the analysis of our model.

Proposition 2.4 (Profitability of Inferior Products). Suppose $\underline{f} > c_n - \bar{a}$ for all n and $v - \underline{f} > w - c_{min}^w$, and consider any proportion κ of sophisticated consumers and any shares s_n, s_n^w . Then, shrouding occurs either with probability one or with probability zero. In the former case, firms sell the superior product w to sophisticated consumers and earn zero profits on it, while they sell the inferior product v to naive consumers and earn positive profits on it. In the latter case, the Bertrand outcome obtains.

Proposition 2.4 says that if $v - \underline{f} > w - c_{min}^w$, then a positive-profit equilibrium in which naive consumers are deceived always exists, and all profits firms earn (in any equilibrium) must derive from selling the *inferior* product. The intuition is in two parts. First, because sophisticated consumers realize that the deceptive product is costly but naive consumers believe it is a better deal, in equilibrium the two types of consumers separate, so—quite in contrast to the message of Section 2.4.1—sophisticated consumers do not create an incentive to unshroud. Second, if a firm unshrouded the additional prices of the inferior product, naive consumers would simply switch to the superior product, so the firm would not attract anyone to its inferior product. Using the logic of our single-product model, the superior product guarantees a deceptive equilibrium with positive profits from the inferior product by rendering the inferior product socially wasteful in comparison.²²

The above insights imply that for socially valuable products, the existence of a superior product can expand the scope of profitable deception and make naive consumers worse off: while unshrouding and marginal-cost pricing can obtain when there is no other product in the market, the opposite happens with the superior product around. Similarly—but perhaps even more perversely—the addition of an *inferior* product can also expand the scope of profitable deception, creating firms’ profit base in (or shifting it to) the new product.

The condition $v - \underline{f} > w - c_{min}^w$, which ensures that a positive-profit deceptive equilibrium exists, 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

²⁰The only exceptions to this rule are sales loads paid to brokers and costs associated with administering separate accounts. These and other regulations are intended to protect the dilution of existing investors’ shares for the benefit of new investors.

²¹In the US market, our model best fits Class B and C mutual fund shares (\$463 billion in total net assets in 2013 according to the ICI fact book 2013), which generally have zero front loads and hence face a binding floor on the up-front price. Class A shares (\$1881 billion in total net assets in 2013) have positive front loads, and are therefore a less obvious fit for our model with a binding price floor. But since front-load fees are charged largely to pay intermediaries’ commissions, intermediaries have an incentive to shroud these fees from investors, so that these fees are arguably also shrouded. Consistent with the view that investors are unaware of these fees, Christoffersen et al. (2013) report that higher front loads indicate higher fund inflows and lower future fund returns. Given that the regulation discussed above implies that other—salient—fees cannot be negative, this means that the salient fees of Class A mutual fund shares also likely face a binding price floor of zero.

²²Of course, as in previous propositions an unshrouded-prices equilibrium also exists. But also as in previous cases, for the reasons outlined in Section 2.2 the deceptive equilibrium is more plausible whenever it exists, and hence we focus our discussion on it.

w . This condition holds if product w is not much better than product v or \underline{f} is not too high. For instance, although a naive consumer may realize that a debit card fulfills the same functions that she wants to use in a credit card, she may still prefer a credit card because she falsely believes that its perks make it a better deal.²³

The market for mutual funds fits this case of our model. Since $c_{min}^w = \underline{f} = 0$ and $v > w$, the sorting condition in Proposition 2.4 is satisfied, so the proposition implies that managed funds will be sold in a deceptive and profitable way. This prediction helps explain the explosion of managed funds, a trend that is often seen as a puzzle due to the existence of superior index funds (Gruber 1996, French 2008): 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.

Note that Proposition 2.4 holds for *any* market shares s_n, s_n^w for the two products. In particular, this means that in our model not even a “specialist” in product w —a firm that sells exclusively or mostly the superior product—has an incentive to unshroud. Intuitively, competition reduces the margin on the superior product to zero, so whether or not it unshrouds a specialist makes no money from the superior product. In Section 2.5.1, we discuss how this result is qualified with market power and costly unshrouding.

Going beyond the setting of our model, the insights above have an immediate implication for the marketing of superior and inferior products: because the inferior product is more profitable, firms have an incentive to push it on consumers who may not otherwise buy it, further decreasing social welfare by expending resources to sell an inferior good. First, firms may pay intermediaries to convince consumers to buy the inferior product. Second, firms may engage in persuasive advertising to induce demand for the inferior product, with the—to the best of our knowledge—novel implication that persuasive advertising is directed exclusively to an inferior good. Third, firms may inform consumers unaware of the inferior product of the product’s existence, yet not do the same for the superior product. Fourth, firms may make costly (real or perceived) improvements to the inferior product to make it more attractive to consumers.²⁴ In some markets, firms may burn all of their gross profits on

²³Although 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 2.4, an equilibrium in which product v is shrouded and product w is unshrouded exists in that case as well. If in addition $w > v$ and there are sufficiently many firms in the market, the *only* profitable equilibrium is the one in which the superior product is unshrouded and the inferior product is shrouded. Consider, for example, a candidate equilibrium in which the superior product w is shrouded. Then, naive consumers must be buying product w ; otherwise, a firm could attract all these naive consumers by setting prices (\underline{f}, \bar{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.

²⁴As an example, many credit cards offer perks, such as airline miles or free rental-car insurance, to attract consumers. The possibility that firms facing a price floor attract consumers by increasing the product’s value rather than decreasing its price raises several interesting considerations that add wrinkles to our results, but that do not seem to affect the logic of our model or fully eliminate the relevance of the price floor. First, in some situations increases in value are subject to a similar adverse-selection-induced constraint as decreases in the price. For instance, if a card offered overly attractive perks, it would attract many sophisticated consumers who are only interested in the perks and not in any of the card’s other services, and who therefore

pushing inferior products, so that they not only sell an overpriced inferior product to naive consumers, they do not even make net profits on it.

It is worth emphasizing that in addition to the overarching main prediction that wasteful and inferior products facilitate profitable deception and hence should go hand in hand with it, our theory has a number of auxiliary empirical predictions when the price floor is binding. First, all firms produce the inferior product, but only the most efficient firms produce the superior product, so (to the extent that there is cost heterogeneity across firms) active firms' market shares tend to be smaller in the market for the inferior product than in the market for the superior product. Relatedly, inferior products tend to go together with less efficient firms, and inferior products attract more entry. Finally, as explained above, firms disproportionately push inferior products. While a full evaluation of the empirical relevance of our theory requires further research, there is some evidence in a number of domains that firms push inferior products. Anagol et al. (2012) and Mullainathan et al. (2010) document that (through intermediaries) firms tend to push inferior products in the life-insurance and mutual-fund markets, respectively. Consistent with the perspective that this hurts consumers, Bergstresser et al. (2009) find that broker-sold funds deliver lower risk-adjusted returns than do direct-sold funds. And Agarwal and Evanoff (2013) document that borrowers eligible for cheaper and superior loans were steered into subprime-like mortgages by brokers or real-estate agents.

2.5 Extensions and Modifications

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

2.5.1 Costly Unshrouding and Market Power

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

Costly unshrouding. Consider the same game as in Section 2.2, except that firm n has to pay $\eta \geq 0$ to unshroud additional prices. Using this modification, we provide one justification for our presumption that firms play a deceptive equilibrium whenever it exists:

avoid the additional price. Second, our framework implies that exactly in an attempt to avoid adverse selection, firms will tie perks to behaviors associated with additional prices. While credit-card issuers clearly do this—e.g., by giving airline miles only for spending on the card, or offering car-rental insurance only if the card is used for the rental—it seems that any such attempt is imperfect. A firm cannot explicitly tie perks to paying additional prices, as this would effectively reveal additional prices to consumers, and to the extent that perks are tied to usage, a sophisticated consumer can use each card only for the specific perk she is interested in, avoiding all or most additional prices. Third, as discussed in Footnote 18, perks are often inefficient, so that they can be interpreted as an inefficient form of pushing rather than as an efficient delivery of service.

Proposition 2.5 (Unique Equilibrium with Costly Unshrouding). Fix all parameters other than η , and suppose that $\underline{f} > c_n - \bar{a}$ for all n and a deceptive equilibrium exists for $\eta = 0$. Then, for any $\eta > 0$ there exists a unique equilibrium, and in this equilibrium all firms shroud and offer (\underline{f}, \bar{a}) with probability one.

Proposition 2.5 says that if a deceptive equilibrium exists in our basic model, then it is the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is. To see the logic of this result, suppose that some firm unshrouds with positive probability. Notice that since $\eta > 0$, in order to unshroud a firm must make positive gross profits afterwards. Hence, no firm unshrouds with probability one—as this would lead to Bertrand-type competition and zero gross profits. Now for each firm that unshrouds with positive probability, take the supremum of the firm’s total price conditional on the firm unshrouding, and consider the highest supremum. At this price, a firm cannot make positive profits if any other firm also unshrouds. Hence, conditional on all other firms shrouding at this price, the firm must make higher profits from unshrouding than from shrouding. But this is impossible: if the firm has an incentive to shroud in this situation with zero unshrouding cost—which is exactly the condition for a deceptive equilibrium to exist—then it strictly prefers to shroud with a positive unshrouding cost.

To draw out further implications of costly unshrouding, we suppose that in addition to being costly, unshrouding reaches and educates only a fraction λ_n of naive consumers. We allow λ_n to differ across firms to capture the notion that some firms (e.g., firms with established marketing departments) may at the same cost be able to reach more consumers than other firms. We assume that if firms n_1, \dots, n_k unshroud, a fraction $\max\{\lambda_{n_1}, \dots, \lambda_{n_k}\}$ of consumers becomes educated. In addition, we suppose that a firm can offer only one pricing scheme (f, a) .²⁵ For simplicity, we only consider the case $\underline{f} + \bar{a} > v$.²⁶

In this setting, the most tempting deviation from a candidate deceptive equilibrium with prices (\underline{f}, \bar{a}) is to unshroud, still charge $f_n = \underline{f}$ to attract a share s_n of the remaining uneducated consumers, and lower a_n to $v - \underline{f}$ to attract all educated consumers. This is unprofitable if the following variant of the Shrouding Condition (SC) holds:

$$s_n(\underline{f} + \bar{a} - c_n) \geq [\lambda_n + (1 - \lambda_n)s_n](v - c_n) - \eta. \quad (2.2)$$

Costly unshrouding somewhat qualifies the distinction between socially wasteful and socially valuable products we have emphasized in Section 2.3.3. By Condition (2.2), it is still true that if the product is socially wasteful, a deceptive equilibrium always exists. It is no longer true, however, that if the product is socially valuable and there are sufficiently many firms, a

²⁵If a firm can offer multiple pricing schemes, from the perspective of deriving conditions for a deceptive equilibrium to exist we can think of educated and uneducated consumers as being in separate markets. Because deviating from deceptive pricing in the market for uneducatable consumers is obviously unprofitable, a firm’s incentives to change its prices and to unshroud derive entirely from the market for educatable consumers. Hence, this case is equivalent to assuming that $\lambda_n = 1$ and firms can offer only one pricing scheme.

²⁶In contrast to our basic model, in a socially valuable industry with an unshrouding cost a deceptive equilibrium might exist if $\underline{f} + \bar{a} < v$; this possibility does not affect the economic conclusions we derive below.

deceptive equilibrium does not exist—even if a firm makes little money on deception and can make positive gross profits transparently selling the product, it may prefer not to undertake costly or partial unshrouding. Nevertheless, a version of our basic message that wastefulness facilitates profitable deception continues to hold: the worse a product is from a social point of view (i.e., the lower is $v - c_n$), the more likely it is that a deceptive equilibrium exists.

While we have not characterized equilibria in general when Condition (2.2) does not hold for all n , an interesting equilibrium arises in the special case in which this condition is violated for a single firm n but

$$(1 - \lambda_n)s_{n'}(\underline{f} + \bar{a} - c_{n'}) \geq \max\{[\lambda_{n'} + (1 - \lambda_{n'})s_{n'}](v - c_{n'}) - \eta, [\lambda_n + (1 - \lambda_n)s_{n'}](v - c_{n'})\} \quad (2.3)$$

for all $n' \neq n$. Then, there is an equilibrium in which firm n educates a fraction λ_n of consumers, while (by Condition (2.3)) all other firms sell to naive consumers. This equilibrium seems consistent with the observation that some firms have attempted to unshroud and sell transparent products to consumers, but they have often had only limited impact on their market. Our theory makes comparative-statics predictions on which firm is most likely to use a transparent strategy. Rewriting the reverse of Condition (2.2) as $\lambda_n(1 - s_n)(v - c_n) - s_n(\underline{f} + \bar{a} - v) - \eta > 0$, we obtain that (ceteris paribus) a firm is more likely to offer a transparent product if (i) it can educate more consumers (λ_n is high); (ii) it, such as a new entrant, has a smaller market share (s_n is low); or (iii) it is more efficient (c_n is lower).²⁷

While it qualifies our results on single-product valuable industries, the addition of an unshrouding cost does not affect the main message of Section 2.4 that there is often an equilibrium in which naive consumers buy a deceptive inferior product: since a firm has no incentive to unshroud the additional prices of the inferior product even if this is free and full, it certainly does not have an incentive to do so if the same thing is costly and partial.

Market power. We now consider the effect of market power of a simple kind—we assume that firm n is strictly more efficient than competitors ($c_n < c_{n'}$ for all $n' \neq n$). We also suppose first that $\eta = 0$. In this case, the condition for when a deceptive equilibrium exists in our single-product model remains unchanged. Intuitively, competitors costs do not affect their prices when shrouding occurs, so these costs do not play a role in determining whether firm n wants to unshroud.

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

Proposition 2.6 (Market Power in the Superior Product). Suppose $\underline{f} > c_{n'} - \bar{a}$ for all n' , $\eta = 0$, and $v - \underline{f} \geq w - c_n^w$.

²⁷If Condition (2.2) is violated for multiple firms, a free-riding issue arises: a firm prefers another firm to pay the unshrouding cost, even if it then prefers to compete in price for educated consumers. This implies that no pure-strategy equilibrium exists, and it is difficult to characterize the equilibrium in general.

I. If $M > s_n(\underline{f} + \bar{a} - c_n)$, there is an equilibrium in which firm n unshrouds the additional prices of product v , other firms shroud, sophisticated consumers and consumers educated by firm n buy product w , and uneducated naive consumers buy product v .

II. If $M \leq s_n(\underline{f} + \bar{a} - c_n)$, there is an equilibrium in which each all firms shroud the additional prices of product v , sophisticated consumers buy product w , and naive consumers buy product v .

Part I of Proposition 2.6 says that if firm n 's market power is larger than its share of deceptive profits, then it prefers to unshroud and attract the naive consumers it can educate. Such a situation could arise, for instance, if firm n sells only the superior product ($s_n = 0$). This observation explains why some specialists in superior products—e.g., Vanguard in the market for mutual funds—have tried to educate consumers about the inferiority of alternative products. At the same time, Part II of Proposition 2.6 says that if firm n 's market power is lower than its share of deceptive profits, then it prefers not to educate consumers. Similarly to Section 2.4.2, the intuition is that deception creates a large profit margin in the inferior product despite competition, and unshrouding the additional prices of this product only leads consumers to buy the less profitable superior one.

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

2.5.2 Inefficiencies in Collecting Revenue from Additional Prices

Our model assumes that the additional price is simply a transfer from consumers to firms. In many settings, however, there may be an inefficiency associated with collecting additional prices. In the printer market, for example, the additional price results from high cartridge prices consumers initially do not understand; but once they do, they reduce their demand in response, generating an inefficiency in cartridge usage. Alternatively, as for instance in the case of credit-card late fees, inefficiencies may obtain from the administrative, customer-service, or legal costs associated with collecting high additional prices.

To capture the inefficiency in charging additional prices, we assume that when a firm chooses the additional price $a \in [0, \bar{a}]$, it imposes a utility loss of a on a consumer and earns profits of $\pi(a)$, where $\pi(0) = 0$, $\pi(a)$ is differentiable, and $\pi'(a) < 1$ for all a . We denote the maximum profits a firm can collect from the additional price by π^* , and let a^* be an additional price that achieves these profits. We begin with characterizing equilibria when the price floor is not binding:

Proposition 2.7 (Non-Binding Price Floor and Collection Inefficiency). Suppose $\underline{f} < c_{min} - \pi^*$.

I. If the product is socially valuable ($v > c_{min}$), unshrouding occurs with probability one

and the Bertrand outcome obtains.

II. If the product is socially wasteful ($v < c_{min}$), then for any $\psi \in [0, 1]$ there exist equilibria in which unshrouding occurs with probability ψ . If shrouding occurs in an equilibrium, consumers buy the product from a most-efficient firm, pay $f = c_{min} - \pi^*$, $a = a^*$, get utility $v - c_{min} - (a^* - \pi^*) < 0$, and firms earn zero profits. If unshrouding occurs in an equilibrium, consumers refrain from buying the product.

An inefficiency in collecting additional prices changes our predictions in one main way: that if the product is socially valuable, only an unshrouded-prices equilibrium exists. Intuitively, if the additional prices were shrouded, firms would—inefficiently—charge the most profitable additional price. A firm could then deviate by unshrouding and eliminating the inefficiency, obtaining positive profits from selling a more efficient product than competitors. This means that even if the price floor is not binding, deception tends to go together with socially wasteful products.

Assuming that additional prices impose an inefficiency does not modify our main qualitative insights regarding when profitable deception occurs in the case of a binding price floor. Analogously to the Shrouding Condition (SC) in our basic model, a deceptive equilibrium exists if

$$s_n(\underline{f} + \pi^* - c_n) \geq v - c_n \text{ for all } n.$$

As in the case of a non-binding price floor, the inefficiency in imposing additional prices increases the welfare costs of deception compared to a situation where $\pi(a^*) = a^*$. Unlike in the case of a non-binding price floor, however, this welfare cost is borne by firms.

2.5.3 Misprediction of Value

Our models in this paper assume that naive consumers mispredict the total price of a shrouded deceptive product. In some settings, it is plausible to assume that instead of or in addition to the price, consumers mispredict the gross value v . For example, a consumer may overestimate the gross return of a managed mutual fund or underestimate the inconvenience associated with claiming an insurance benefit. In this section, we discuss how such mispredictions of value affect our conclusions. We assume that the true value of the product is v and the perceived value is $\tilde{v} > v$, and that unshrouding eliminates consumers' misperception regarding value. We consider two cases depending on whether unshrouding also eliminates consumers' misperception regarding the total price.

First, we suppose that unshrouding eliminates misperceptions regarding value, but not regarding price. This would be the case, for instance, for education that sheds light on the quality of customer service in the insurance industry. Then, since lowering consumers' perceived valuation of a homogenous product cannot increase a firm's demand, in our single-product model no firm ever has an incentive to unshroud, so a deceptive equilibrium always exists. In our multi-product model, in any deceptive equilibrium naive consumers perceive the value of the inferior product to be \tilde{v} , so the sorting condition in Proposition 2.4 must be rewritten as $w - c_{min}^w < \tilde{v} - \underline{f}$ —that the inferior product generate sufficiently high surplus taking consumers' misperceived value into account. But if this condition holds, a deceptive

equilibrium exists. Intuitively, since unshrouding cannot increase a firm’s demand for the inferior product, and—because the superior product has a profit margin of zero—the firm cannot make money from selling the superior product, there is no incentive to unshroud.

Second, we suppose that unshrouding eliminates consumers’ misperceptions regarding both the value and the price. This captures, for instance, education that explains the net return of managed funds relative to index funds to consumers. In our single-product model, the only situation in which consumers’ value enters a firm’s considerations is when unshrouding occurs, and in this situation consumers have correct perceptions of the product’s value. Hence, a firm contemplating deviation from a candidate deceptive equilibrium faces exactly the same incentives as in our basic model, so the same Shrouding Condition determines whether a deceptive equilibrium exists. The logic of our multi-product model, instead, is modified in a way similar to that above. In any deceptive equilibrium, naive consumers perceive the value of the inferior product to be \tilde{v} , so the sorting condition in Proposition 2.4 becomes $w - c_{min}^w < \tilde{v} - \underline{f}$. If this condition holds, a deceptive equilibrium exists for the same reason as in Proposition 2.4: given that a superior product is available at marginal cost, no firm can make money once consumers perfectly understand the nature of the products, so no firm has an incentive to unshroud.

2.5.4 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 2.3, a deceptive equilibrium with prices (\underline{f}, \bar{a}) often exists because unshrouding would lead consumers with values between \underline{f} and $\underline{f} + \bar{a}$ not to buy, discretely reducing industry demand. The deceptive equilibrium is more likely to exist when there are more such consumers—that is, when there are more consumers who are mistakenly buying the product. And a deceptive equilibrium exists whenever the product could not be profitably sold to consumers who understand its total price. As above, this is the case whenever the product is socially wasteful to produce, for example because no consumer values it above marginal cost, or (in a natural extension of our model) the number of such consumers is insufficient given some fixed costs of production. But if the product can be profitably sold in a transparent way, then with a sufficient number of firms at least one firm would choose to unshroud, eliminating the deceptive equilibrium.

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

unshrouding.²⁸

2.6 Related Theoretical Literature

In this section, we discuss theories most closely related to our paper. Relative to the existing literature, our main contribution is identifying the central role of socially wasteful and inferior products in profitable deception, and drawing out additional implications of this finding.

In Gabaix and Laibson’s (2006) model, firms sell a base good with a transparent price and an add-on with a shroudable price, and consumers buying the base good can avoid the add-on by undertaking costly steps in advance. Gabaix and Laibson’s main prediction is that unshrouding the add-on prices can be unattractive because it turns profitable naive consumers (who fail to avoid the expensive add-on) into unprofitable sophisticated consumers (who avoid the add-on). Although the precise trade-off determining a firm’s decision of whether to unshroud is different, we start from a similar insight, and use it to explore a number of new implications.²⁹

Clarifying and adapting Gabaix and Laibson’s theory, Armstrong and Vickers (2012) investigate a model of contingent charges in financial services, and apply it the UK retail banking industry. Consistent with our perspective, they argue that the “free if in credit” model—whereby firms charge nothing for account maintenance, and rely on contingent charges, such as overdraft protection, for revenue—can be naturally explained by the presence of naive consumers.³⁰

In research complementary to ours, Grubb (2012) considers services, such as cellphone calls or bank-account 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

²⁸The 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 \bar{a}_n ; in fact, our proof of Propositions 2.3 and 2.5 in the Appendix allows for this possibility.

²⁹Beyond its main message regarding the role of wasteful and inferior products in maintaining profitable deception, our theory identifies some other economically important forces that Gabaix and Laibson’s model abstracts from. For instance, we show that because in a multi-product market naive and sophisticated consumers might separate, sophisticated consumers do not necessarily make markets more transparent. And we establish that if imposing additional prices is inefficient, then unshrouding is more likely to occur.

³⁰Relatedly, Piccione and Spiegler (2012) characterize how firms’ ability to change the comparability of prices through “frames” affects profits in Bertrand-type competition. If a firm can make products fully comparable no matter what the other firm does—which is akin to unshrouding in our model and that of Gabaix and Laibson (2006)—the usual zero-profit outcome obtains. Otherwise, profits are positive. Piccione and Spiegler highlight that increasing the comparability of products under any frame through policy intervention will often induce firms to change their frames, which can decrease comparability, increase profits, and decrease consumer welfare. Investigating different forms of government interventions, Ko and Williams (2011) and Kosfeld and Schüwer (2011) demonstrate that educating naive consumers in the Gabaix and Laibson’s (2006) framework can decrease welfare because formerly naive consumers may engage in inefficient substitution of the add-on.

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.³¹

Our theory is also related to the classical industrial-organization literature on markets in which firms sell a primary product as well as a complementary good or service, consumers are locked in after the primary purchase, and consumers do not observe a firm’s price for the complementary good. The main goal of this literature is to analyze the effect of market power in the market for the complementary good on prices and consumer welfare. Closely related to our benchmark Proposition 2.1, Lal and Matutes (1994) show that in competing for consumers, sellers offset the high price for the complementary good with a low price for the primary good. Nevertheless, Shapiro (1995), Hall (1997), and Borenstein, Mackie-Mason and Netz (2000) show that social welfare may not be maximized due to deadweight losses from consumers’ reactions to the overly low price of the primary good and the overly high price of the complementary good. Our paper has a different objective—analyzing the scope for and implications of profitable deception in a competitive market—and shows that deception can lead to additional welfare losses from the systematic sale and pushing of inferior products.

Relatedly, in assuming that consumers can be induced to pay high additional fees once they buy a product, our theory shares a basic premise with the large literature on switching costs. But even if firms cannot commit to ex-post prices and there is a floor on ex-ante prices—so that positive profits obtain in equilibrium—our model’s main insights do not carry over to rational switching-cost models. Most importantly, we are unaware of any rational switching-cost model that predicts the systematic sale of inferior products in competitive markets, and under the natural assumption that consumers know or learn product attributes, it is clear that a firm is indeed better off selling a superior product.³²

Our result that firms sell profitable inferior products to unknowing consumers may also seem reminiscent of a similar potential implication in classical pricing models based on asymmetric information. But the fact that in our model consumers mispredict prices leads to important differences. For instance, while in a rational asymmetric-information model a lower-quality inferior product may be more profitable to sell than a superior product because it is cheaper to produce and consumers do not know its value, in our theory a lower-quality inferior product may be more profitable than a superior product even if it is more expensive

³¹Our theory also builds on a growing literature in behavioral industrial organization that assumes consumers are not fully attentive, mispredict some aspects of products, or do not fully understand their own behavior. See for instance DellaVigna and Malmendier (2004), Eliaz and Spiegel (2006), Spiegel (2006a, 2006b), Laibson and Yariv (2007), Grubb (2009), and Heidhues and Kőszegi (2010).

³²This point is immediate if consumers know the products’ values at the time of original purchase. As an illustration of the same point when consumers learn product values only after initial purchase, suppose that there are two products with values v_H and v_L and costs c_H and c_L , respectively, product H is strictly superior ($v_H - c_H > v_L - c_L$), and the switching cost is k . Consider a firm who could sell either product to a consumer at an ex-ante price of p followed by an ex-post price of its choice. If the firm sells the superior product, it can charge an ex-post price of $c_H + k$ in a competitive market, so it makes profits of $p + c_H + k - c_H$. If the firm sells the inferior product, then (as the consumer realizes that she can switch to the superior product) it can charge an ex-post price of $c_H + k - (v_H - v_L)$, so it makes profits of $p + c_H + k - (v_H - v_L) - c_L$. It is easy to check that the firm prefers to sell the superior product.

to produce and consumers do know its value. More importantly, in a rational model the sale of a profitable inferior product cannot be systematic: while an uninformed rational consumer may in *some* states buy a product that is inferior to and more profitable than an alternative available product, this cannot be the case *on average* for the states in which she buys the product.^{33,34}

2.7 Some Policy Implications and Conclusion

Beyond the main theme of the paper that we have developed above—that in equilibrium socially wasteful and inferior products might be sold profitably to consumers—a further important general message emerges from our results: that deception is likely to be more widespread and economically harmful in the case of a binding than in the case of a non-binding price floor. First, while in a deceptive equilibrium with a non-binding price floor competitive pricing and productive efficiency obtain—all consumers buy from a most-efficient firm and pay a total price equal to the firm’s cost—in a deceptive equilibrium with a binding price floor neither does—consumers pay a higher price and also buy from less efficient firms. Second, the profits firms earn from deception when the price floor is binding can potentially attract entry of less efficient producers, exacerbating productive inefficiency. Third, in a deceptive equilibrium with a binding price floor firms push inferior products, and to the extent that such pushing is costly and expands the market share of inferior relative to superior products, it causes social harm. Fourth, when the price floor is not binding—but not when it is binding—the existence of a deceptive equilibrium for socially valuable products is fragile to our assumptions of perfect symmetry between firms and the efficiency of deception. For instance, if one firm has slightly lower marginal cost but the other can impose a higher additional price, then the former firm has an incentive to unshroud to gain a competitive edge. And as we show in Section 2.5, when there is an inefficiency in collecting additional prices, then for socially valuable products a deceptive equilibrium cannot be maintained with a non-binding price floor, but might be maintained with a binding price floor.

The suboptimal nature of market equilibrium in our models raises the obvious question of whether a social planner can improve outcomes. We discuss here a few issues related to potential policy interventions, but a fuller exploration of this difficult question is outside the

³³To see this, suppose that a consumer buys product v at price p in states S , and that there is a product w available at price p^w in these same states. Supposing for simplicity that all firms have the same marginal costs of production, denote the cost of products v and w by c and c^w , respectively. If product v is inferior to product w conditional on S —i.e., $E[w - c^w | S] > E[v - c | S]$ —yet product v is also more profitable conditional on S —i.e., $p^w - E[c^w | S] < p - E[c | S]$ —then $E[w | S] - p^w > E[v | S] - p$, so that the consumer would be better off buying product w in the states in which she buys product v .

³⁴Similarly, when consumers must pay classical search costs to find out prices (or product features), at a broad level one can think of the prices as being partly shrouded. While we believe that search costs are very important in the markets we consider, by themselves they do not seem to fully explain why inferior or socially wasteful products are systematically sold in a more profitable way than superior products. Furthermore, although we have no precise empirical evidence, it does not seem that firms are playing the mixed-strategy pricing equilibrium predicted by these models.

scope of this paper.³⁵

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

An alternative, and in the context of deception more direct, way of intervening in the market is to decrease \bar{a} by regulating hidden charges. For example, the Credit Card Accountability, Responsibility, and Disclosure (Credit CARD) Act of 2009 limits late-payment, over-the-limit, and other fees to be “reasonable and proportional to” the consumer’s omission or violation, thereby preventing credit-card companies from using these fees as sources of extraordinary ex-post profits. Similarly, in July 2008 the Federal Reserve Board amended Regulation Z (implementation of the Truth in Lending Act) to severely restrict the use of prepayment penalties for high-interest-rate mortgages. Regulations that require firms to include all non-optional price components in the up-front price—akin to recent regulations of European low-cost airlines—can also serve to decrease \bar{a} .

To see the effects of such regulations in our setting, consider a policy that decreases the maximum additional price firms can charge from \bar{a} to $\bar{a}' < \bar{a}$ in the range where the price floor is binding. If Condition (SC) still holds with \bar{a} replaced by \bar{a}' , firms charge $(\underline{f}, \bar{a}')$ in the new situation, so that the decrease in the additional price benefits consumers one to one. And if the decrease in the additional price leads to some firm violating Condition (SC), the market becomes transparent, prices drop further, and productive efficiency obtains. The prediction that consumers may benefit from controls on one component of the price that appears fully fungible with another component provides both a counterexample to a central argument brought up against many consumer-protection regulations—that its costs to firms will be passed on to consumers—and a testable prediction of our model when the price floor is binding. Consistent with this prediction, Bar-Gill and Bubb (2012) and Agarwal, Chomsisengphet, Mahoney and Stroebel (2013) document that the Credit CARD Act—while succeeding in lowering regulated fees—did not lead to an increase in unregulated fees or a decrease in the availability of credit, so that it lowered the total cost of credit to consumers.

It is important to note, however, that the direct regulation of additional prices is unlikely to be effective as a general approach to combating deception. As Agarwal et al. (2009) discuss in their setting, formulating ex-ante guidelines for which charges are acceptable seems extremely difficult, and individually researching and approving each new financial product is very costly. One potential response to this challenge is for researchers to develop portable empirical methods for detecting additional prices and consumer mistakes from commonly available data. As a simple example, the empirical approach of Chetty, Looney and Kroft (2009) takes advantage of the general observation that if a consumer reacts to changes in

³⁵See Congdon, Kling and Mullainathan (2011) and especially Agarwal, Driscoll, Gabaix and Laibson (2009) for further discussion of the advantages and disadvantages of different types of policies aimed at mitigating the effects of consumer mistakes.

financially identical price components differently, then she is not paying sufficient (relative) attention to some components of the price. Another response to the same challenge is to develop policies that do not rely on the social planner knowing which price components consumers underappreciate. For example, Murooka (2013) shows that when financial intermediaries motivated by commissions are used by firms to sell deceptive products, capping commissions, or requiring commissions to be uniform across products, can improve welfare even if the regulator cannot identify additional prices.

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

Finally, in this paper we have taken the opportunity to deceive consumers—that is, the shrouded additional price component—as exogenous. In most real-world markets, however, someone has to come up with ways to hide prices from consumers, so that the search for deception opportunities can be thought of as a form of innovation. In our companion paper (Heidhues et al. 2012b), we study the incentives for such “exploitative innovation,” and contrast them with the incentives for innovation that benefits consumers. Here too we find a perverse incentive: because learning ways to charge consumers higher additional prices increases the profits from shrouding and thereby lowers the motive to unshroud, a firm may have a strong incentive to make exploitative innovations *and* have competitors copy them. 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. In addition, we show that if the deceptive feature can be copied by competitors, the incentive to come up with deceptive products is stronger when the price floor is binding than when it is not. Since deception often seems based on contract features that can be easily copied, it is therefore likely to be more widespread when the price floor is binding, adding another reason that deception is a greater problem in this case than when the price floor is not binding. The possibility of exploitative innovation also adds caution to policies that directly regulate the maximum additional price: such a policy can greatly increase firms’ incentive to make new exploitative innovations, and hence may have a small net effect.

2.8 Proofs

Proof of Proposition 2.1. First, we establish that for any $\psi \in [0, 1]$ there exists an equilibrium with the properties stated in the proposition. Let every firm offer the contract $f_n = c_n - \bar{a}$, $a_n = \bar{a}$ and unshroud with probability γ , where $(1 - \gamma)^N = 1 - \psi$. In this equilibrium, every firm makes zero profits independent of whether shrouding occurs or not,

and—again independent of whether shrouding occurs or not—no firm can attract customers at a total price above its marginal cost. In what follows, we establish that any equilibrium has the properties specified in the proposition.

First, suppose unshrouding occurs with probability $\psi = 1$. Then all consumers buy from a firm that charges the lowest total price $f_n + a_n$ as long as this total price is below v . We therefore have Bertrand competition in the total price, and standard arguments imply that the Bertrand outcome obtains.

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

Second, no firm makes sales at a total price $f_n + a_n > c_{min}$ with positive probability. To see this, suppose otherwise. Then, some firm makes sales with positive probability at a total price above c_{min} . Any most-efficient firm, i.e. a firm with cost $c_n = c_{min}$, can copy this firm's strategy and thereby earn positive profits. Hence, any most-efficient firm makes positive profits in equilibrium. Denote by $\pi_n^* > 0$ a most-efficient firm n 's equilibrium profits. Let \bar{f} be the supremum of up-front prices of any most-efficient firm conditional on shrouding occurring. We consider two cases: (i) $\bar{f} > \underline{f}$ and (ii) $\bar{f} = \underline{f}$.

Consider case (i). Suppose first that some most-efficient firm n sets the up-front price \bar{f} with positive probability. Denote by \hat{a} the supremum of firm n 's additional price conditional on shrouding and setting an up-front price of \bar{f} . Then, only firm n can set \bar{f} with positive probability conditional on shrouding, as otherwise firm n would benefit from minimally undercutting the up-front price. Hence, conditional on charging (\bar{f}, a_n) and shrouding, firm n can earn positive expected profits only when unshrouding occurs and all rivals charge a total price equal to or above $\bar{f} + a_n$ with positive probability. Thus, with positive probability all rivals must charge total prices weakly above $\bar{f} + \hat{a}$, and charge total prices strictly above $\bar{f} + \hat{a}$ if firm n charges $\bar{f} + \hat{a}$ with positive probability.

Let \bar{t}_n be the supremum of the total prices of firm n conditional on unshrouding. Let $\bar{t} = \max\{\bar{t}_n | c_n = c_{min}\}$. Consider a most-efficient firm, say firm n' , that achieves \bar{t} when unshrouding. First, suppose firm n' sets \bar{t} when unshrouding with positive probability. Because no firm other than n' can set a total price \bar{t} with positive probability, firm n' must earn positive profits when all other most-efficient firms shroud. This contradicts that $\bar{t} > \bar{f} + \hat{a}$. Thus, no firm sets \bar{t} with positive probability. Note that $\bar{t} > \bar{f} + \hat{a}$ if firm n charges $\bar{f} + \hat{a}$ with positive probability and $\bar{t} \geq \bar{f} + \hat{a}$ otherwise. Consider a sequence $t_{n'} \rightarrow \bar{t}$ of optimal total prices by firm n' when unshrouding, and denote the corresponding expected total profits by $\epsilon_{n'}$. Then, $\epsilon_{n'} \rightarrow 0$ contradicting the fact that n' must earn strictly positive profits.

Suppose, thus, that no most-efficient firm charges \bar{f} with positive probability when shrouding. Let firm n be a firm that achieves this supremum. Note that firm n 's profits conditional on shrouding occurring go to zero as $f_n \rightarrow \bar{f}$. Consider a sequence of optimal offers (f_n^l, a_n^l) when shrouding for which $f_n^l \rightarrow \bar{f}$. This sequence must have a convergent subsequence in which (f_n^l, a_n^l) converges to (\bar{f}, \hat{a}) . From now on consider this subsequence

and note that the corresponding profits conditional on shrouding occurring go to zero. Hence for sufficiently high l , firm n must make profits conditional on unshrouding occurring. Thus, $\bar{t} \geq \bar{f} + \hat{a}$. Furthermore, $\bar{t} > \bar{f} + \hat{a}$ if firm n charges $\bar{f} + \hat{a}$ with positive probability because no other most-efficient firm can with positive probability set a total price equal to $\bar{f} + \hat{a}$ when unshrouding in this case. From here on, the proof proceeds as in the above case where some firm sets \bar{f} with positive probability.

Consider case (ii). Since $\bar{f} = \underline{f}$, any most-efficient firm sets \underline{f} with probability one conditional on shrouding. Furthermore, because $\underline{f} + \bar{a} \leq c_{min}$ and no firm sets a price below c_{min} with positive probability, any most-efficient firm must set $a_n = \bar{a}$ conditional on shrouding. Thus, conditional on shrouding occurring, all most-efficient firms make zero profits, and hence they must make zero profits in equilibrium. Furthermore, no other firm than a most-efficient firm can sell at a total price strictly above c_{min} because then a most-efficient firm could make positive profits by mimicking that firm.

We have thus established that no firm sells a contract at a total price other than c_{min} . Next, observe that consumers must buy the product at a total price of c_{min} with probability one if $v > c_{min}$; otherwise, there exists $\epsilon > 0$ for which a most-efficient firm can profitably deviate by unshrouding and charging a total price $c_{min} + \epsilon$ with probability one. Furthermore, conditional on unshrouding occurring, if $v < c_{min}$ consumers buy the product with probability zero because no firm makes sales at a total price strictly below c_{min} with positive probability.

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

Proof of Proposition 2.2. With consumers who observe and take into account the additional price, we have Bertrand competition in the total price and the proposition follows from standard arguments. \square

Proof of Proposition 2.3. We establish a slightly more general version of this proposition: we allow the maximum additional prices firms can impose to differ across firms. Let \bar{a}_n be the maximum additional price firm n can impose. We prove the statement of Proposition 2.3 with Inequality (SC) replaced by

$$s_n(\underline{f} + \bar{a}_n - c_n) \geq v - c_n. \quad (2.4)$$

Note first that if some firm unshrouds with probability one, all other firms are indifferent between shrouding and unshrouding. Thus, an equilibrium in which unshrouding occurs with probability one always exists. In any such unshrouded-prices equilibrium consumers observe and take the additional price into account, so that our game reduces to a standard

Bertrand game in which the consumers' willingness to pay is v . Hence, in any unshrouded-prices equilibrium consumers buy the product only if $v \geq c_{min}$. In case $v > c_{min}$, standard Bertrand-competition arguments imply that all consumers buy the product at total price of c_{min} from a most-efficient firm.

Now consider deceptive equilibria, i.e., equilibria in which shrouding occurs with probability one. In case firm n has a positive probability of sales in equilibrium, it must set $a_n = \bar{a}_n$ as otherwise it could increase its profits conditional on a sale by increasing a_n without affecting the probability of selling. The same argument as in the text establishes that if Inequality (2.4) holds for all n , then there is a deceptive equilibrium in which all firms set $(\underline{f}, \bar{a}_n)$. We now provide a formal argument for why firms set \underline{f} with probability one in any deceptive equilibrium. The proof is akin to a standard Bertrand-competition argument. Take as given that all firms shroud with probability 1, and that firm n sets the additional price \bar{a}_n . Note that by setting $f_n = \underline{f}$, firm n can guarantee itself a profit of $s_n(\underline{f} + \bar{a}_n - c_n) > 0$. As a result, no firm will set $f_n > v$, because then no consumer would buy from it. Take the supremum \bar{f} of the union of the supports of firms' up-front price distributions. We consider two cases. First, suppose that some firm sets \bar{f} with positive probability. In this case, all firms have to set \bar{f} with positive probability; otherwise, a firm setting \bar{f} would have zero market share and hence zero profits with probability one. Then, we must have $\bar{f} = \underline{f}$; otherwise, a firm could profitably deviate by moving the probability mass to a slightly lower price. Second, suppose that no firm sets \bar{f} with positive probability. Let firm n 's price distribution achieve the supremum \bar{f} . Then, as f_n approaches \bar{f} , firm n 's expected market share and hence expected profit approaches zero—a contradiction.

We now establish that if the strict inequality

$$s_n(\underline{f} + \bar{a}_n - c_n) > v - c_n \quad (2.5)$$

holds for all n , in any equilibrium shrouding occurs either with probability one or zero. Suppose otherwise. Conditional on all competitors shrouding with positive probability, firm n can guarantee itself positive profits by shrouding itself and offering the contract $(\underline{f}, \bar{a}_n)$, which attracts consumers since $v \geq \underline{f}$ and makes positive profits since $\underline{f} + \bar{a}_n > c_n$. To earn positive profits when unshrouding, firm n must set a total price $t_n \leq v$. Consider the supremum of the total price \hat{t}_n set by firm n when unshrouding, and let $\hat{t} = \max_n \{\hat{t}_n\}$. Note that there exists at most one firm that sets this price with positive probability; if two or more firms did so, then some firm could increase profits by moving this probability mass to slightly below \hat{t} . First, consider the case in which one firm puts positive probability mass on \hat{t} , and let this firm be firm n . Firm n earns zero profits conditional on some other firm unshrouding. Conditional on all others shrouding it earns less than

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

and hence would increase its profits by deviating from unshrouding and charging the total price \hat{t} to shrouding and charging $(\underline{f}, \bar{a}_n)$. Second, consider the case in which no firm puts positive probability mass on \hat{t} . Let n be a firm that achieves this supremum. Consider a sequence $t_n \rightarrow \hat{t}$ of optimal total prices by firm n when unshrouding, and denote the

corresponding expected profits conditional on some other firm unshrouding by ϵ_n . Then, $\epsilon_n \rightarrow 0$. Notice that conditional on all other firms shrouding, firm n earns at most $t_n - c_n \leq v - c_n$, so that the (unconditional) expected profit of firm n is less than $v - c_n + \epsilon$, which by Condition (2.5) is strictly less than $s_n(\underline{f} + \bar{a}_n - c_n)$ for a sufficiently small $\epsilon > 0$. Hence, firm n is strictly better off shrouding and charging $(\underline{f}, \bar{a}_n)$ —a contradiction. Thus, if Condition (2.5) holds, unshrouding occurs with probability one or zero in equilibrium.

In the remainder of this proof, we establish by contradiction that if Condition (2.4) is violated for some firm, then in any equilibrium additional prices are unshrouded with probability one. Note again that if unshrouding occurs with probability one, we have Bertrand competition in the total price, and hence the Bertrand outcome obtains. The proof that unshrouding occurs with probability one in this case proceeds in three steps.

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

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

First, suppose all firms other than n set a total price strictly higher than \hat{t} with positive probability. Because each firm $n' \neq n$ makes zero profits when unshrouding occurs, it must make positive profits when shrouding occurs. In addition, since it only makes profits when shrouding occurs, it sets the additional price $\bar{a}_{n'}$ with probability one. Take the supremum of firms' up-front prices \bar{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), $\bar{f}' \leq v$. Note that $\bar{f}' + \bar{a}_{n'} > \hat{t}$ for any $n' \neq n$.

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

If only one firm n' sets \bar{f}' with positive probability, then firm n' has zero market share both when unshrouding occurs and when shrouding occurs and some firm other than n' sets a total price strictly greater than \hat{t} . Because firm n' earns positive profits by (i) and is the only firm that sets \bar{f}' with positive probability conditional on the total price being strictly higher than \hat{t} , every firm except for n' sets its up-front price strictly higher than \bar{f}' and its total price weakly lower than \hat{t} when shrouding with positive probability. Suppose first $n' = n$. Then, there exists a firm $n'' \neq n$ that shrouds and sets an up-front fee $f_{n''} > \bar{f}'$, $a_{n''} \leq \hat{t} - f_{n''}$ with positive probability. Since $\bar{f}' + \bar{a}_m > \hat{t}$ for any $m \neq n$, $\bar{a}_{n''} > \hat{t} - \bar{f}'$. Then,

firm n'' can increase its profits by decreasing all prices $f_{n''} > \bar{f}'$ to \bar{f}' and by increasing its additional price holding the total price constant—a contradiction. Next, suppose $n' \neq n$. Then, firm n shrouds, sets $f_n > \bar{f}'$ with positive probability and charges an additional price $a_n \leq \hat{t} - f_n$ with probability one when charging these up-front prices. For almost all of these up-front prices, firm n must earn strictly positive profits when shrouding occurs; otherwise firm n could unshroud with probability one and guarantee positive profits when all rivals shroud and charge a total price above \hat{t} . Thus, firm n' shrouds and sets $f'_n \geq f_n > \bar{f}'$, $a_{n'} \leq \hat{t} - f_{n'}$ with positive probability. Since $\bar{f}' + \bar{a}_m > \hat{t}$ for any $m \neq n$, firm n' can increase its profits by decreasing all prices $f_{n'} > \bar{f}'$ to \bar{f}' and increasing its additional price holding the total price constant—a contradiction.

If no firm sets \bar{f}' with positive probability, there exists firm n' that for any $\epsilon > 0$ sets up-front prices in the interval $(\bar{f}' - \epsilon, \bar{f}')$ with positive probability. As $\epsilon \rightarrow 0$, the probability of firm n' 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. Thus, we establish that $\bar{f}' = \underline{f}$.

Because $\bar{f}' = \underline{f}$, 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, $\underline{f} + \bar{a}_{n'} \geq \hat{t}$ for any $n' \neq n$. We now show that whenever shrouding, any firm $n' \neq n$ does not set 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 $g' > \underline{f}$ such that firm n' sets prices above g' with positive probability. There cannot be a competitor whose up-front price when shrouding falls on the interval $[\underline{f}, g']$ with positive probability; otherwise, firm n' could increase its profits by decreasing all prices above g' to \underline{f} and by increasing its additional price holding the total price constant. But then, firm n' can raise its up-front price from \underline{f} to 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 than or equal to \hat{t} when shrouding. But if it unshrouds and sets the same prices, it would also earn profits when all rivals shroud and set a price above \hat{t} , thereby strictly increasing its profits—a contradiction. Hence firm n also must set \underline{f} with probability one when shrouding.

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

Step (iii): Additional prices are unshrouded with probability one. Suppose not. Then,

each firm chooses to shroud with positive probability. Take the infimum of total prices \underline{t} set by any firm when shrouding. We consider two cases. First, suppose $\underline{t} \leq v$. Take a firm that achieves the infimum. By (i), this firm earns positive profits. For any $\epsilon > 0$, take total prices below $\underline{t} + \epsilon$ of the firm. By unshrouding and setting $\underline{t} - \epsilon$, the firm decreases its profits by at most 2ϵ when one or more other firms unshroud, but discretely increases its market share if all other firms shroud. Hence, for sufficiently small $\epsilon > 0$ this is a profitable deviation—a contradiction. Second, suppose $\underline{t} > v$. Take firm n that violates Inequality (2.4). 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 unshrouds, and its profits are at most $s_n(\underline{f} + \bar{a}_n - c_n)$ when shrouding occurs. But then, deviating and setting a total price equal to v is profitable because conditional on others shrouding firm n would earn $v - c_n > s_n(\underline{f} + \bar{a}_n - c_n)$. \square

Proof of Proposition 2.4. First, if unshrouding occurs with probability one, all consumers are sophisticated and standard (Bertrand-competition) arguments imply that they buy a total-surplus-maximizing product at marginal cost.

From now on, suppose shrouding occurs with positive probability. Our proof has four steps. In Step (i) we show that if shrouding occurs with positive probability, firms earn positive profits. Step (ii), which is contained in Lemma 2.1, establishes that every firm sets an up-front price $f_n = \underline{f}$ when shrouding. Since $v - \underline{f} > w - c_{min}^w$, Step (ii) implies that naive consumers buy the inferior product whenever shrouding occurs. Step (iii) uses Bertrand-competition-type arguments to show that sophisticated consumers always buy a total-surplus-maximizing product at marginal cost. This implies (Step (iv)) that firms can only earn profits if shrouding occurs, and hence all firms shroud with probability one and sell the inferior product at prices (\underline{f}, \bar{a}) .

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

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

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

Lemma 2.1. $\tilde{d} = v - \underline{f}$.

Because $v - \underline{f} > w - c_{min}$ and no firm sells the superior product below cost with pos-

itive probability in equilibrium, Lemma 2.1 implies that when shrouding occurs all naive consumers buy the inferior product at an up-front price equal to the floor.

Proof of Lemma 2.1. Suppose, toward a contradiction, that $\tilde{d} \neq v - \underline{f}$. Since no firm can set an up-front price below the floor, $\tilde{d} < v - \underline{f}$. We begin by showing that there is a firm n' such that for any $\epsilon > 0$, firm n' has an optimal action that involves shrouding and giving perceived deal $\tilde{d}_{n'}(\epsilon)$ such that $0 \leq \tilde{d}_{n'}(\epsilon) - \tilde{d} < \epsilon$ and the expected profits of firm n' when choosing this action and shrouding occurring is less than ϵ . This is trivial from the definition of \tilde{d} if no firm sets \tilde{d} with positive probability when shrouding, or exactly one firm sets \tilde{d} with positive probability when shrouding. Suppose, therefore, that two or more firms set \tilde{d} with positive probability when shrouding. Let n' be one of these firms. We will show that if firm n' earns positive expected profits from this action conditional on shrouding occurring, then it has a profitable deviation. If it earns positive expected profits on the inferior product only, lowering $f_{n'}$ minimally while still shrouding and holding other prices constant is a profitable deviation; if it earns positive expected profits on both the inferior and the superior product, then lowering $f_{n'}$ and $t_{n'}^w$ minimally by the same amount is a profitable deviation; and if it earns positive expected profits on the superior product only, then unshrouding and lowering $t_{n'}^w$ minimally is a profitable deviation. This concludes our argument that a firm n' with the above properties exists. Denote a deal offered from a firm n to sophisticated consumers, or its “real deal,” by $d_n = \max\{w - t_n^w, v - f_n - a_n\}$. For future reference, note two further facts: (i) our argument implies that if firm n' chooses \tilde{d} and shrouds, conditional on shrouding occurring it earns zero profits; and (ii) we can choose sequences $\epsilon^l \rightarrow 0$ and $\tilde{d}_{n'}(\epsilon^l)$ such that the corresponding real deals converge; denote the limit by $d(\tilde{d})$.

For a sufficiently small $\epsilon > 0$, the optimal action in which firm n' sets the perceived deal $\tilde{d}_{n'}(\epsilon)$ does not earn its equilibrium expected profits conditional on shrouding occurring, so firm n' must earn $\pi_{n'}^* - \epsilon$ conditional on unshrouding occurring. Now for it to be the case that firm n' makes $\pi_{n'}^* - \epsilon$ when unshrouding occurs but at most ϵ when shrouding occurs, there must exist a firm $n'' \neq n'$ and an $\epsilon' > 0$ such that firm n'' offers a strictly better real deal than the real deal associated with $\tilde{d}_{n'}(\epsilon)$ with probability of at least $1 - \epsilon'$ when it shrouds but offers a worse real deal than that associated with $\tilde{d}_{n'}(\epsilon)$ with positive probability when it unshrouds. Furthermore, as $\epsilon \rightarrow 0$, we must have $\epsilon' \rightarrow 0$, and the probability with which firm n'' offers a worse real deal than that associated with $\tilde{d}_{n'}(\epsilon)$ remains bounded away from zero. This implies that there is a firm n'' such that firm n'' offers a weakly better real deal than $d(\tilde{d})$ with probability 1 when it shrouds but offers a weakly worse real deal than $d(\tilde{d})$ with positive probability when it unshrouds.

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

First suppose that some firm offers \underline{d} with positive probability when unshrouding. Observe that two firms cannot offer this deal with positive probability, as otherwise a firm would

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

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

$$\begin{aligned} v - f' &> \max\{w - \hat{t}^w, v - \hat{f}\}, \\ v - f' - a' &< \max\{w - \hat{t}^w, v - \hat{f} - \hat{a}\}, \end{aligned}$$

where (f', a') is a contract offer that corresponds to the perceived deal \tilde{d}' . Since firm n' must attract naive consumers with positive probability when offering \tilde{d}' , it is strictly better off changing the perceived deal \hat{d} to \tilde{d}' while holding the real deal constant. To complete the argument, we show that this deviation is feasible. In case $f' + a' \geq \hat{f} + \hat{a}$, firm n' can do so by lowering \hat{f} to f' and raising the additional price from a' so as to keep the total price of the inferior product fixed. If $f' + a' < \hat{f} + \hat{a}$, then the second inequality above implies

that $\max\{w - \hat{t}^w, v - \hat{f} - \hat{a}\} = w - \hat{t}^w$ and hence firm n' can do so by changing (\hat{f}, \hat{a}) to (f', a') , thereby improving the perceived deal without affecting the real deal $w - \hat{t}^w$ it offers to sophisticated consumers as well as naive consumers when unshrouding occurs. We thus have a contradiction.

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

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

Hence, we conclude that $\tilde{d} = v - \underline{f}$, and therefore all firms set an up-front price \underline{f} for the inferior product with probability one conditional on shrouding and all naive consumers

buy the inferior product when shrouding occurs with probability one since $v - \underline{f} > w - c_{min}^w$ and no firm in equilibrium sells the superior product at a price $t_n^w < c_{min}^w$ with positive probability.

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

Step (iv): Since no firm can offer more than the maximal total surplus $w - c_{min}^w$ without making a loss, and only a most-efficient firm can make an offer in which $d_n = w - c_{min}^w$ without making a loss, with probability one all sophisticated consumers buy from a most-efficient firm. Furthermore, because no firm can earn positive profits if unshrouding occurs, every firm shrouds with probability one. And to earn positive profits, every firm must set a real deal for the inferior product that is strictly worse than $w - c_{min}^w$. Hence, only naive consumers buy the inferior product. Because every firm must attract some naive consumers to earn positive profits, every firm shrouds with probability one and sets $f_n = \underline{f}$, and, since firm n 's market share is independent of a_n , every firm sets $a_n = \bar{a}$. This complete the proof of the proposition. \square

Proof of Proposition 2.5. Again, as in the proof of Proposition 2.3 we establish a slightly more general version of this proposition in which the maximum additional prices firms can impose differ across firms. Let \bar{a}_n be the maximum additional price firm n can impose. We prove the statement of Proposition 2.5 maintaining the assumption that the price floor is

binding for all firms, i.e., $\underline{f} > c_n - \bar{a}_n$ for all n . Recall from the proof of Proposition 2.3 that with different maximal additional prices, a deceptive equilibrium exists for $\eta = 0$ if and only if for all firms n ,

$$s_n(\underline{f} + \bar{a}_n - c_n) \geq v - c_n. \quad (2.6)$$

This proof has five steps.

Step (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 2.2, all consumers buy at a total price $f + a = c_{min}$ and no firm makes positive profits from selling to the consumers excluding the unshrouded cost. Then, the firm that chooses to unshroud makes negative profits—a contradiction.

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

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

Step (iv): No firm unshrouds the additional price with positive probability. Let \hat{t}_n be the supremum of the equilibrium total-price distribution of firm n when unshrouding; set $\hat{t}_n = 0$ in case firm n does not unshroud. Let $\hat{t} = \max_n \{\hat{t}_n\}$; by (iii), \hat{t} is bounded from above and hence well-defined. Consider firm n that unshrouds and for whom $\hat{t}_n = \hat{t}$. Note that in any equilibrium in which some firm unshrouds with positive probability, $\hat{t} > c_n$ by (ii).

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

Second, suppose that firm n does not charge the total price \hat{t} with positive probability. Then, for any $\epsilon > 0$, firm n charges a total price in the interval $(\hat{t} - \epsilon, \hat{t})$ with positive probability. As $\epsilon \rightarrow 0$, the probability that firm n conditional on some other firm unshrouding can attract consumers goes to zero, because \hat{t} is the supremum of the total-price distribution conditional on unshrouding. Hence, firm n cannot earn the unshrouding cost $\eta > 0$ conditional on some other firm unshrouding—i.e., it loses money in expectation relative to shrouding and offering $(\underline{f}, \bar{a}_n)$. In addition, conditional on all other firms shrouding firm n

earns less than the deviation profits in the no-unshrouding-cost case. Because shrouding is an equilibrium in the no-unshrouding-cost case, there is a profitable deviation for firm n —a contradiction.

Step (v): All firms offer the contract $(\underline{f}, \bar{a}_n)$ with probability one. By (iv), all firms choose to shroud with probability one. Hence, in equilibrium all firms charge an additional price $a = \bar{a}_n$ with probability one. By the exact same argument as in Proposition 1.3, all firms offer up-front price $f = \underline{f}$. \square

Proof of Proposition 2.6. Part I. Consider a candidate equilibrium in which all firms other than n charge marginal cost in the superior-product market, firm n charges c_{min}^w in this market, all firms charge (\underline{f}, \bar{a}) in the inferior-product market, and only firm n unshrouds. According to the tie-breaking assumption for the superior-product market, therefore, firm n attracts all superior-product consumers. Since lowering a attracts no new customers in the inferior-product market and raising f induces all inferior-product customers to buy from a rival, charging (\underline{f}, \bar{a}) in the inferior product market is optimal for all firms. Furthermore, no firm $n' \neq n$ can profitably attract customers in the superior-product market, and thus both charging marginal cost and shrouding is a best response. It is also clearly optimal for firm n to charge c_{min}^w in the superior-product market, because in order to attract uneducated naive consumers firm n has to set a price of product w below its production cost. Furthermore, the decision of whether to shroud does not affect the behavior of sophisticated consumers, and hence we can focus on the profits earned from naive consumers. When shrouding, firm n earns

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

while when unshrouding it earns

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

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

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

Proof of Proposition 2.7. The proof closely follows that of Proposition 2.1, and we only sketch the differences.

Since now $\pi(a_n) < a_n$ for all $a_n > 0$, if unshrouding occurs with probability one ($\psi = 1$) standard arguments imply that the Bertrand outcome obtains in which $a_n = 0$.

As in Proposition 2.1, no firm sells the product at a price at which $f_n + \pi(a_n) < c_{min}$ with positive probability. Suppose that $\psi < 1$ and some firm sells the product at $f_n + \pi(a_n) > c_{min}$ with positive probability. Then, any most-efficient firm can make positive profits. Now following the exact same arguments as those in proof of Proposition 2.1 leads to a contradiction. Hence, in equilibrium $f_n + \pi(a_n) = c_{min}$ and again by the arguments in the proof of Proposition 2.1, consumer choose a contract $f_n = c_{min} - \pi(a^*), a_n = a^*$ if shrouding occurs. If the industry is socially valuable ($v > c_{min}$), however, a most-efficient

firm could unshroud, offer the contract $f_n = c_{min} + (a^* - \pi(a^*))/2$, $a_n = 0$, and thereby ensure positive profits. Hence, if the industry is socially valuable, unshrouding occurs with probability one. Next, if the industry is socially wasteful ($v < c_{min}$), no firm can benefit from unshrouding and hence there exists an equilibrium in which all most-efficient firms offer $f_n = c_{min} - \pi(a^*)$, $a_n = a^*$ and thereby earn zero profits, inefficient firms choose a higher upfront price and total price, and one firm unshrouds with probability $\psi \in [0, 1]$ while all other firms shroud. \square

Chapter 3

Exploitative Innovation (with Paul Heidhues and Botond Köszegi)

3.1 Introduction

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

In this paper, we analyze the incentives for exploitative innovation in a market for potentially deceptive products, and contrast them with the often-studied incentives for making product improvements consumers value. Section 3.2 introduces our model, which consists of a simultaneous-move price-competition stage modeled after Gabaix and Laibson (2006) and Heidhues et al. (2012a), and a preceding innovation stage. At the price-competition stage, firms selling perfect substitutes each set a transparent up-front price as well as an additional price, and unless at least one firm decides to (costlessly) unshroud additional prices, naive consumers ignore these prices when making purchase decisions. To capture the notion that in some markets, such as banking services, credit cards, and mutual funds, firms cannot return all profits from later charges by lowering initial charges—even if they can dissipate those profits in other ways—we deviate from most existing work and posit that there is a floor on the up-front price.¹ We assume that whenever a deceptive equilibrium—wherein

¹As we discuss in Section 3.2, in Heidhues et al. (2012) we provide a microfoundation for the price floor based on the presence of “arbitrageurs” who would take advantage of overly low prices. This microfoundation is an extreme variant of Ellison’s (2005) insight that firms may be reluctant to cut initial prices because these cuts disproportionately attract less profitable consumers. While we assume that firms may not be able to return all profits by lowering the up-front price, however, our model is consistent with the possibility that

all firms shroud additional prices—exists at the pricing stage, firms play that continuation equilibrium.²

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

In Section 3.3, we characterize innovation incentives when the price floor is not binding—a condition we argue holds for some commonly invoked examples of deceptive products, including hotel rooms and printers—and find that they are, identically for exploitative and value-increasing innovations, based on the “appropriable part” of the innovation. If firm 1 increases its product value by \$10 relative to others, it can charge almost \$10 more than competitors and still capture the entire market, earning a profit of almost \$10 per consumer. And if firm 1 figures out a way to charge a \$10 higher additional price than others, it can charge slightly lower prices than competitors, capture the entire market, and make \$10 more ex post, again earning a profit of almost \$10 per consumer. This can help explain why firms have developed some appropriable exploitative practices, such as a proprietary technology that prevents printer users from buying non-brand printer cartridges, in industries with a non-binding price floor. But because firms have no incentives to make non-appropriable innovations, and the primary tools for exploitation often seem to be easily copyable contract terms, the extent of deception in these industries may in the end be limited.

In Section 3.4, we characterize innovation incentives when the price floor is binding—a condition we argue holds for a number of consumer financial products, including credit cards and banking services—and find that they are now often stronger for exploitative than for value-increasing innovation. Our results can be understood in terms of firm 1’s incentive to maintain a deceptive equilibrium in the face of the threat that a competitor—unable to compete on the up-front price—prefers to unshroud and compete on the additional price. This deviation allows a competitor to capture the entire market, but because consumers who learn of the additional prices may not buy the product, it may require a discrete cut in the total price. A profitable deceptive continuation equilibrium, therefore, exists if and only if each firm’s profit when sharing deceived consumers at a high profit margin is greater than its profit when unshrouding and capturing all consumers at a low profit margin.

The above logic implies that so long as a profitable deceptive continuation equilibrium exists, firm 1 has strong incentives for both appropriable and non-appropriable exploitative innovation. Since in a deceptive equilibrium the floor on the up-front price is binding,

firms dissipate profits in other ways—such as through entry costs or advertising—and hence earn zero net profits. We discuss this issue in more detail in the conclusion.

²Whenever such an equilibrium exists, it is the most plausible one for two reasons. Most importantly, it is then the unique equilibrium in the variant of our model in which unshrouding is costly, no matter how small the cost is. In addition, in most of our settings a deceptive equilibrium Pareto dominates an unshrouded-prices equilibrium from the perspective of the firms, and in some settings it does so strictly.

firm 1 always benefits from an exploitative innovation by being able to increase its margin. Furthermore, because acquiring the exploitative innovation increases competitors' profits from shrouding and thereby lowers their motive to unshroud, the innovation may enable profitable deception in the industry. In this case, firm 1 *prefers* competitors to acquire the exploitative innovation—so that it prefers non-appropriable to appropriable innovations—and its willingness to pay is equal to its full post-innovation profits. In contrast, firm 1 can never benefit from a non-appropriable value-increasing innovation, and because such an innovation increases competitors' motive to unshroud by raising the profits from selling a transparent product, a firm may even be willing to pay to *avoid* the innovation.

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

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

Our paper builds on a theoretical literature in behavioral economics that explores how firms use hidden fees or otherwise take advantage of mistakes in consumer decisions.³ To our knowledge, however, neither this literature, nor the extensive classical literature on research and development, has investigated firms' incentives to make exploitative innovations. In fact, our message that a firm may be eager to invest in socially wasteful non-appropriable

³See, for instance, DellaVigna and Malmendier (2004), Eliaz and Spiegler (2006), Spiegler (2006a, 2006b), Grubb (2009, 2012), Heidhues and Köszegi (2010), Piccione and Spiegler (2012), and Ko and Williams (2011), and Spiegler (2011) for a textbook treatment.

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

3.2 Basic Model

In this section, we introduce our model of innovation in a market for potentially deceptive products. The game has two stages, an innovation stage and a price-setting stage. The price-setting stage is a variant of the model in Heidhues et al. (2012a), and we begin by describing this stage. There are $N \geq 3$ firms competing for a unit mass of naive consumers who value firm n 's product at $v_n > 0$ and are looking to buy at most one item. Firms simultaneously set up-front prices f_n and additional prices a_n , and decide whether to costlessly unshroud the additional prices. The highest possible additional price firm n can charge is $\bar{a}_n > 0$, and a consumer buying product n has to pay both prices f_n and a_n . If all firms shroud, consumers make purchase decisions believing that the total price of product n is f_n . If at least one firm unshrouds, all consumers become aware of all additional prices and hence make purchase decisions based on the total prices $f_n + a_n$.⁴

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

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

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

Each firm’s cost of providing the product is $c > 0$. We assume that $v_n + \bar{a}_n > c$ for all n ; a firm with $v_n + \bar{a}_n < c$ cannot profitably sell its product, so without loss of generality we can think of it as not participating in the market. We also impose that $\underline{f} \leq v_n$ for all n , so that in a shrouded market consumers are willing to buy from a firm with an up-front price at the floor. Finally, we make two simple tie-breaking assumptions. First, consumers go to a highest-quality firm when indifferent.⁶ Second, if all firms shroud and a subset of firms with the same quality choose an up-front price at the floor, these firms split their demand in proportion to shares $s_n \in [0, 1)$.

We now turn to describing the innovation stage. To identify innovation incentives in a transparent manner, we assume that only one firm, firm 1, can make innovation investments. Firms start from a symmetric position in v_n and \bar{a}_n , with $v_n = v$ and $\bar{a}_n = \bar{a}$ for all n . Then, firm 1 chooses whether or not to invest in innovation, with all firms observing its decision. We consider separately two types of innovation. An “exploitative innovation” costs I_a and increases the maximum additional price firm 1 can charge by Δa and the maximum additional price firm $n \neq 1$ can charge by $\Delta a'$, where $0 \leq \Delta a' \leq \Delta a$. This formulation allows us to consider the two extreme cases often studied in the literature, appropriable innovations (which competitors cannot copy: $\Delta a' = 0$) and non-appropriable innovations (which competitors can fully copy: $\Delta a' = \Delta a$), as well as in-between cases. Analogously, a value-increasing innovation costs I_v and increases consumers’ valuation of product 1 by Δv and their valuation of product $n \neq 1$ by $\Delta v'$, where $0 \leq \Delta v' \leq \Delta v$.⁷

We look for subgame-perfect Nash equilibria in the game played between firms, imposing—as is standard in the industrial-organization literature—that no firm charges a total price below its marginal cost. In addition, we assume that if a deceptive Nash equilibrium—wherein all firms shroud additional prices—exists in the pricing subgame, firms play that continuation equilibrium. Because no firm has an incentive to shroud if at least one firm unshrouds, there is always an unshrouded-prices continuation equilibrium. When a deceptive continuation equilibrium exists, however, it is more plausible than the unshrouded-prices continuation equilibrium for two main reasons. Most importantly, in any situation in which we use this assumption, the deceptive equilibrium is then the unique equilibrium in the variant of our model in which unshrouding carries a positive cost, no matter how small the cost is.⁸ In addition, in most of our settings a deceptive equilibrium Pareto dominates an

⁶This allows a firm to price a lower-quality competitor out of the market by offering the same deal—something it could do anyway by offering a minimally better deal—ensuring the existence of a pure-strategy equilibrium. This simplifies some of our proofs, but does not affect the logic of our results.

⁷While we interpret value-increasing innovations as increasing the product’s true value to consumers, the same results hold for innovations, advertisement, and other investments that merely increase the *perceived* value—with the investment’s social value of course being lower in this case than for true value-increasing innovations. For example, a mutual-fund prospectus outlining an investment philosophy may fool consumers into believing that there is a dependable way to beat the market, increasing the perceived value of the fund.

⁸The proof of Proposition 5 in Heidhues et al. (2012a) establishes this claim for any situation in which consumers value firms’ products equally and the price floor is binding, a condition that applies to Propositions 3.2 and 3.3 below. For Part II of Proposition 3.1 and Proposition 3.4 below, we establish the claim in Appendix B. In Proposition 3.5 below, we use our equilibrium refinement only for the subgame following no innovation, and in this case consumers value products equally and the price floor is binding. In Part I of

unshrouded-prices equilibrium from the perspective of the firms, and in some settings it does so strictly. Using these refinements, we characterize investment incentives by identifying the maximum investment costs I_a^* and I_v^* below which firm 1 is willing to make the investment of each type.

3.3 Non-Binding Price Floor

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

Proposition 3.1 identifies innovation incentives in this case:

Proposition 3.1. Suppose $\underline{f} \leq c - (\bar{a} + \Delta a)$ for all n . Then,

- I. (Value-Increasing Innovation). $I_v^* = \Delta v - \Delta v'$.
- II. (Exploitative Innovation). $I_a^* = \Delta a - \Delta a'$.

Part I of Proposition 3.1 says that firm 1's incentive to make a value-increasing innovation is based on the “appropriable part” of the innovation—the extent to which the innovation increases the value of its product above that of competitors' products. As a simple example, suppose firms' cost is \$100, innovation increases the value of firm 1's product from \$200 to \$220 and that of the other firms from \$200 to \$210, and the maximum additional price—which firms actually charge in any deceptive continuation equilibrium—is \$50. Similarly to the logic of Lal and Matutes (1994), classical switching-cost models, and many existing behavioral-economics models with naive consumers, firms compete aggressively for ex-post-profitable consumers, and bid down the up-front price to \$100-\$50=\$50. Absent the innovation, therefore, firm 1 cannot sell its product above an up-front price of \$50, so it earns zero net profits. If it innovates, however, firm 1 can charge an up-front price slightly below \$60 and attract all consumers, generating total revenue of nearly \$110 with the additional price included. Hence, firm 1 earns a profit of \$10 per consumer.

Part II of Proposition 3.1 says that similarly to its incentive to make a value-increasing innovation, firm 1's incentive to make an exploitative innovation is equal to the appropriable part of the innovation. In this case, however, firm 1's competitive advantage derives not from offering a better product to consumers, but from better exploiting consumers. Continuing with the above example, suppose the innovation increases firm 1's maximum additional price to \$70 and other firms' maximum additional price to \$60. Because of the higher additional price they can charge, other firms are now willing to bid down the up-front price to \$40.

Proposition 3.1, we do not use the refinement.

Even so, firm 1 can offer a slightly lower up-front price, and again earn a revenue of nearly \$110 with its higher additional price included. And because other firms are charging a total price equal to marginal cost, no firm can profitably unshroud and attract consumers. In other words, the profitability of exploitative and value-increasing innovations is exactly the same: a value-increasing innovation allows a firm to raise its total price above competitors' and still keep consumers, and an exploitative innovation allows a firm to lower its price below competitors' and still make profits.

An example consistent with the prediction that a firm will make appropriable exploitative innovations is the printer industry. Hall (1997) describes a number of strategies, including questionable “artistic” cartridge design patents, printer-head patents, and perpetual design modifications, that generate no consumer value but help printer manufacturers control the cartridge market and thereby cash in on naive consumers. Nevertheless, with a non-binding price floor a firm’s incentive to make exploitative innovations is no greater than its incentive to make value-increasing innovations, and in particular it has no incentive to make non-appropriable exploitative innovations. Because the primary tools for deception often seem to be contract terms that tie the consumer to the firm and induce her to pay supra-normal fees ex post, and such contract innovations are typically easy to copy, the extent of deception in industries with a non-binding price floor may in the end be limited.

3.4 Binding Price Floor

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

To analyze innovation incentives, we first identify conditions under which a deceptive equilibrium exists in the pricing subgame, assuming that the maximum additional price firm n can charge is \bar{a}_n and the value of its product is v . The analysis mirrors that in

⁹If credit-card companies handed out substantial sums to new consumers (i.e., if the up-front price was substantially negative), then many deal-seeking consumers would likely get—and then not use—these cards just to earn the handouts. The threat of losing money on these “arbitrageur” consumers imposes a floor on the up-front price that is likely not much below \$0. Evans and Schmalensee (2005) estimate that the *average* cost of opening a new account, including all marketing and processing cost, is about \$72. And as argued for instance by Ausubel (1991), credit-card companies make large ex-post profits on charges consumers do not anticipate.

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

Heidhues et al. (2012a). Note that if additional prices remain shrouded, all firms set their maximum additional price \bar{a}_n . Then, since firms are making positive profits and hence have an incentive to attract consumers, they bid down the up-front price to \underline{f} . With consumers being indifferent between firms, firm n gets market share s_n and therefore earns a profit of $s_n(\underline{f} + \bar{a}_n - c)$. For this to be an equilibrium, no firm should want to unshroud additional prices and undercut competitors. Once a firm unshrouds, consumers will be willing to pay exactly v for its product, so that firm n can make profits of at most $v - c$ by unshrouding and capturing the entire market. Hence, unshrouding is unprofitable for firm n if

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

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

Lemma 3.1 (Equilibrium in the Pricing Subgame). Suppose $\underline{f} > c - \bar{a}_n$ for all n . If Inequality (SC) holds for all n , a deceptive continuation equilibrium exists. In any deceptive continuation equilibrium, $f_n = \underline{f}$ and $a_n = \bar{a}_n$ for all n , and firms earn positive profits. If Inequality (SC) is violated for some n , in any continuation equilibrium prices are unshrouded with probability one, consumers buy at a total price of c , and firms earn zero profits.

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

In Heidhues et al. (2012a), we show that profitable deception is likely to be more pervasive in socially wasteful industries—where the value consumers derive from the product is below marginal cost—than in socially valuable industries—where the consumer value is strictly above marginal cost. If the product is socially wasteful, a firm that unshrouds cannot profitably sell its product, so—with no firm ever wanting to unshroud—a profitable deceptive equilibrium always exists. But if the product is socially valuable, a firm that would make sufficiently low profits from deception can earn higher profits from unshrouding and capturing the entire market, so if there is such a firm only a non-deceptive, zero-profit equilibrium exists. We also show that the relevant notion of social wastefulness is relative to the best alternative: whether or not the product yields positive social surplus, if an alternative product with higher

social surplus is available, there is typically a deceptive equilibrium in which naive consumers buy the inferior product and firms earn positive profits from it.

Our key results derive from considering how an innovation by firm 1 affects firms' willingness to go along with deceiving consumers, as captured by the Shrouding Condition (SC). Proposition 3.2 states our results for non-appropriable innovations, showing that they are much stronger for exploitative than for value-increasing innovations:

Proposition 3.2 (Non-Appropriable Innovations). Suppose $\underline{f} > c - \bar{a}$. Then,

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

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

Part I of Proposition 3.2 says that if firms can maintain a deceptive equilibrium following the innovation, firm 1 is willing to spend resources on—socially clearly wasteful—non-appropriable exploitative innovation. In this situation, increasing the additional price from \bar{a} to $\bar{a} + \Delta a$ cannot lead to a decrease in the up-front price, so the innovation increases firm 1's profits by at least $s_1 \Delta a$. Going further, if in addition firms cannot maintain a deceptive equilibrium without the innovation, the innovation increases firm 1's profits even more by enabling profitable deception in the industry. In this case, firm 1's willingness to pay for the exploitative innovation is equal to its full post-innovation profits—a potentially huge incentive to innovate.

Part II of Proposition 3.2 says that in contrast to socially-wasteful exploitative innovation, firm 1 is not willing to spend on—socially often beneficial—non-appropriable value-increasing innovation. Because such innovation can increase neither one's market share nor one's markup, firm 1 has no incentive to invest in it. Furthermore, it may be the case that firms can maintain a deceptive equilibrium without but not with the innovation, so that firm 1's willingness to pay for the innovation is negative. Intuitively, an increase in v does not affect profits when firms shroud but—by increasing the amount consumers are willing to pay for a transparent product—does increase the profits a firm can gain from unshrouding. As a result, firm 1 may be willing to spend money to avoid an increase in v .

To sharpen the intuitions from Proposition 3.2 on a firm's incentive to invest in exploitative innovation, we compare the incentives for non-appropriable and appropriable innovations:

Proposition 3.3 (Appropriability of Exploitative Innovation). Suppose $\underline{f} > c - \bar{a}$. Then, I_a^* is weakly greater if the innovation is non-appropriable ($\Delta a = \Delta a'$) than if it is only partially copyable ($\Delta a \geq \Delta a'$), and it is strictly greater if the Shrouding Condition holds for all n when $\bar{a}_n = \bar{a} + \Delta a$ but fails for some $n \neq 1$ when $\bar{a}_n = \bar{a} + \Delta a'$.

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

competitors who are not very good at imposing additional prices gain little from deception and hence may want to deviate from it, threatening the deceptive equilibrium and thereby firm 1’s profits. To eliminate such a threat, firm 1 would like to teach competitors how to better exploit consumers.¹¹

The message of Proposition 3.2 that firms might be willing to make investments in non-appropriable innovations, and that such innovations are likely to be exploitative rather than socially valuable, seems consistent with how some consumer financial products have developed recently. As has been argued by many researchers, consumers likely underappreciate a number of future fees and other payments associated with credit cards, bank accounts, and non-traditional mortgages, and firms exacerbate these mispredictions with carefully designed contract features—such as teaser rates, high fees for certain patterns of product use, and difficult-to-understand payment schedules involving large future payments—whose main purpose is likely to hide products’ total price from consumers. These exploitative contract innovations seem easy to copy, and in fact have been quickly copied by competitors.

Furthermore, not only are the above types of innovations easy to copy, in some instances firms seem—consistent with Proposition 3.3—positively willing to share them with each other. 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.”¹² For a participating issuer, any innovation is essentially a non-appropriable innovation. Proposition 3.2 explains why a participating firm makes innovations, and although there may be other considerations, Proposition 3.3 provides a strategic reason for why a firm that is interested in developing new exploitative practices is willing to join Argus in the first place.

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

Proposition 3.4 (Value-Increasing Innovation in Socially Wasteful Industries). Suppose $\underline{f} > c - \bar{a}$, $v + \Delta v < c$, and $\Delta v > \Delta v'$. Then, $I_v^* = [(1 - s_1)(\underline{f} + \bar{a} - c)] + [\Delta v - \Delta v']$.

Proposition 3.4 implies that firm 1’s willingness to pay for fully or partially appropriable value-increasing innovations in a socially wasteful industry—that is, for products that should not be in the market in the first place—is quite high: it is greater than in the corresponding classical setting (where it would clearly be zero), it is greater than the increase in the relative value of firm 1’s product ($\Delta v - \Delta v'$), and (because I_v^* is bounded away from zero) it is non-trivial even for vanishingly small product improvements. Firm 1’s willingness to pay, I_v^* ,

¹¹An important caveat to Proposition 3.3 is that while firm 1 is willing to share an exploitative innovation with competitors already in the marketplace, it often prefers not to share the same innovation with potential entrants. The ability to charge higher additional prices can make participation in the market more attractive, inducing additional entry and thereby reducing the innovator’s market share.

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

derives from two sources. First, as captured in the first term, the innovation attracts the consumers of all competitors to firm 1, and firm 1 benefits from this even at pre-innovation market prices. Second, as captured in the second term, because the innovation improves firm 1's product more than competitors' products, firm 1 can increase the up-front price without losing consumers, further increasing its profits. Although firm 1 makes these extra profits by pricing competitors out of the market, with the industry being socially wasteful competitors do not unshroud in response.

We next consider socially valuable industries:

Proposition 3.5 (Value-Increasing Innovation in Socially Valuable Industries). Suppose $\underline{f} > c - \bar{a}$, $v > c$, $\Delta v > \Delta v'$, and (SC) holds for all n when $\bar{a}_n = \bar{a}$. Then, $I_v^* = [\Delta v - \Delta v'] - \underline{s}_1(\underline{f} + \bar{a} - c)$.

Proposition 3.5 implies that firm 1's willingness to pay for partially appropriable value-increasing innovations in a socially valuable deceptive industry is quite small: it is lower than in the corresponding classical setting (where it would equal $\Delta v - \Delta v'$), it is lower than in a socially wasteful industry with the same relative product values, and it is negative for non-substantial improvements. In this industry, any partially appropriable innovation must lead to unshrouding; otherwise, firm 1 would be able to price competitors out of the market while setting a high total price, and competitors could unshroud and profitably undercut this total price. Hence, innovation leads firm 1 to lose its positive profits from deception. This loss dampens firm 1's incentive to innovate, and for small improvements—which generate only a small relative advantage—the incentive is negative.

3.5 Concluding Remarks

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

If opportunities to dissipate profits arise *after* the innovation stage, however, our results may be affected. In particular, if firms can offer perks at the pricing stage that are just as effective in attracting consumers as decreases in the up-front price, our theory becomes equivalent to one without a price floor. But it seems plausible to assume that non-price forms of competition are often less effective in attracting consumers, and are subject to decreasing

returns. For example, while mail solicitations clearly generate some demand for credit cards, the returns to sending multiple solicitations to a consumer within a short time frame likely diminish quickly. In this case, firms do not compete away all profits from the additional price at the pricing stage. As a result, the logic of our main insights remains unchanged, albeit with an important qualification to our prediction that a firm always prefers a non-appropriable to an appropriable exploitative innovation. Specifically, a firm is still willing to spend on a non-appropriable innovation, and prefers it over an appropriable innovation if the innovation is pivotal in preventing another firm from unshrouding. But because a fraction of the extra profits from the innovation is competed away, if the innovation is not pivotal in preventing unshrouding the firm prefers to appropriate it.

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

3.6 Proofs and Supplementary Materials

3.6.1 Proofs

Proof of Proposition 3.1.

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

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

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

Let \hat{t}_k be the supremum of the total price distribution firm $k \neq 1$ charges; and let \hat{t}_1 be that of firm 1. Define the quality adjusted maximum of these suprema as $\hat{t} = \max\{\hat{t}_1 - \Delta v, \hat{t}_k - \Delta v'\}$. Note that firm k cannot charge this quality-adjusted total price with positive probability when unshrouding; if it did, it would lose to firm 1 with probability one, contradicting the fact that it must earn positive profits with any price it charges with positive probability. Furthermore, firm k cannot charge this quality-adjusted price with positive probability when shrouding. If it did, it must have positive market share and it can do so only if all other firms shroud. But then it must set $a_k = \bar{a}$ to maximize profits, and hence it offers a contract $(\hat{t} + \Delta v' - \bar{a}, \bar{a})$ with positive probability. For this contract to have a positive market share, firm 1 must shroud and set base prices above $\hat{t} + \Delta v - \bar{a}$ 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 $\bar{a} - \Delta v$. But this is not a best response: firm 1 could keep the total price distribution fixed but always charge the maximal additional price \bar{a} . Then its base price would always lie (weakly) below $\hat{t} + \Delta v - \bar{a}$, 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 $\hat{t} + \Delta v' - \bar{a}$ 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 \bar{a} , thereby ensuring that its base price is below $\hat{t} + \Delta v' - \bar{a}$, 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 $\epsilon \rightarrow 0$, 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 $\hat{t} + \Delta v' - \bar{a}$ while setting total prices strictly below $\hat{t} + \Delta v'$. Firm k in this case could keep its total price distribution fixed, and always charge the maximal additional price \bar{a} , increasing its market share for a set of total prices it charges with positive probability. The argument for why the market share of firm k must go to zero when shrouding and charging quality-adjusted total prices in $(\hat{t} - \epsilon, \hat{t})$ is analogous. We conclude that firm 1 earns its relative advantage in every equilibrium.

Part II. It is easy to check that the following constitutes an equilibrium in the pricing subgame: all firms shroud and set $f_n = c - (\bar{a} + \Delta a')$ and their maximum additional price, firm 1 gets the entire market, and earns a profit of $\Delta a - \Delta a'$. We show that profits are the same in any continuation equilibrium in which firms shroud with probability one, immediately implying the proposition. Clearly, in any continuation equilibrium in which firms shroud with

probability one, all firms set the maximum additional price. Then, the game is equivalent to a Bertrand-competition game in which firm 1 has cost $c - (\bar{a} + \Delta a)$ and all other firms have cost $c - (\bar{a} + \Delta a')$. It is well-known that in any equilibrium of this game when no firm charges below marginal cost, firm 1 earns a profit equal to $\Delta a - \Delta a'$. \square

Proof of Lemma 3.1. See proof of Proposition 3 in Heidhues et al. (2012a). \square

Proof of Proposition 3.2. We first prove Case I. In the subgame following an innovation by firm 1, the Shrouding Condition holds for all firms by assumption, and thus firm 1 earns $s_1(\underline{f} + \bar{a} + \Delta a - c)$ in this case. In the subgame in which firm 1 did not innovate, firm 1 earns $s_1(\underline{f} + \bar{a} - c)$ if the Shrouding Condition holds for all firms and zero otherwise. In the former case the innovation increases firm 1's profits by $s_1\Delta a$, in the latter case by $s_1(\underline{f} + \bar{a} + \Delta a - c)$, which is strictly greater than $s_1\Delta a$ because $\underline{f} + \bar{a} > c$.

We now prove 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 $s_1(\underline{f} + \bar{a} - c)$ which is positive and independent of v . The result, hence, follows from the fact that an increase in v either does not affect whether the Shrouding Condition holds or leads to a violation of the Shrouding Condition for some firm. \square

Proof of Proposition 3.3. Firm 1 earns $s_1(\underline{f} + \bar{a} + \Delta a - c)$ in the subgame following its innovation if the Shrouding Condition holds for all firms and zero profits otherwise. An increase in $\Delta a'$ for some $n \neq 1$ increases the left-hand-side of the Shrouding Condition and hence relaxes it; thus it either does not affect firm 1's profits or—if it makes the Shrouding Condition hold for some firm $n \neq 1$ for which it does not hold otherwise—strictly increases firm 1's profits. \square

Proof of Proposition 3.4.

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

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

To complete the proof, we show that in any pricing subgame following innovation in which firms shroud with probability 1, firm 1 charges prices $\underline{f} + \Delta v - \Delta v', \bar{a}$ with probability one and gets the entire market, so that our equilibrium-selection criterion selects such an

equilibrium. This means that firm 1 earns $\underline{f} + \Delta v - \Delta v' + \bar{a} - c$ if it innovates and $s_1(\underline{f} + \bar{a} - c)$ if it does not. I_v^* is the difference between these two profit levels.

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

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

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

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

Proof of Proposition 3.5. Absent innovation, firm 1 earns $s_1(\underline{f} + \bar{a} - c)$ in the deceptive equilibrium we select. The proof of Part I of Proposition 3.1, which applies unaltered when there is a binding price floor, establishes that firm 1 earns $\Delta v - \Delta v'$ in the pricing subgame following a value-increasing innovation. Thus, $I_v^* = \Delta v - \Delta v' - s_1(\underline{f} + \bar{a} - c)$. Since $\underline{f} + \bar{a} > c$, this cut-off is strictly less than that in Proposition 3.4. \square

3.6.2 Uniqueness of Deceptive Equilibrium with Unshrouding Cost

Part II of Proposition 3.1. We show that if firms face an unshrouding cost $\eta > 0$, then in any continuation equilibrium shrouding occurs with probability 1. The proof is by contradiction and has three steps.

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

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

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

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

Now, let \tilde{f} be the supremum of the up-front prices set by firms other than j conditional on shrouding. Then, for firm j to earn its its equilibrium expected profits with (\hat{f}, \hat{a}) , it must be that $\tilde{f} \geq \hat{f}$. We first rule out equality. Suppose, toward a contradiction, that $\tilde{f} = \hat{f}$. If a

firm other than j sets \tilde{f} with positive probability, firm j prefers to undercut; if not, firm j earns zero profits, in either case a contradiction. Hence, $\tilde{f} > \hat{f}$.

Let \tilde{f}'' , \tilde{a}'' be some prices set by firms other than j when shrouding that satisfy $\tilde{f}'' > \hat{f}$, and suppose it is firm k that sets these prices. Let π_k be firm k 's equilibrium expected profit. By an argument structurally identical to that in the first paragraph of step (iii), for \tilde{f}'' sufficiently close to \tilde{f} firm k cannot earn a profit of π_k when shrouding occurs, so that it must earn positive expected profits when unshrouding occurs. Hence, $\tilde{f}'' + \tilde{a}'' \leq \hat{t}$.

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

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

Proposition 3.4. Obvious: an unshrouding firm makes zero gross profits as it cannot profitably sell the product, so that it makes negative net profits. \square

Bibliography

- Agarwal, Sumit and Douglas D. Evanoff**, “Loan Product Steering in Mortgage Markets,” 2013. Working Paper.
- , **John C. Driscoll, Xavier Gabaix, and David Laibson**, “The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation,” *Brookings Papers on Economic Activity*, 2009, pp. 51–117.
- , **Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel**, “Regulating Consumer Financial Products: Evidence from Credit Cards,” 2013. Working Paper.
- Anagol, Santosh and Hugh Hoikwang Kim**, “The Impact of Shrouded Fees: Evidence from a Natural Experiment in the Indian Mutual Funds Market,” *American Economic Review*, 2012, 102 (1), 576–593.
- , **Shawn Cole, and Shayak Sarkar**, “Understanding the Incentives of Commissions Motivated Agents: Theory and Evidence from the Indian Life Insurance Market,” 2012. Working Paper.
- Armstrong, Mark and John Vickers**, “Consumer Protection and Contingent Charges,” *Journal of Economic Literature*, 2012, 50 (2), 477–493.
- Asker, John and Heski Bar-Isaac**, “Raising Retailers’ Profits: On Vertical Practices and the Exclusion of Rivals,” *American Economic Review*, 2013. Forthcoming.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger**, “Working Over Time: Dynamic Inconsistency in Real Effort Tasks,” 2013. Working Paper.
- Ausubel, Lawrence M.**, “The Failure of Competition in the Credit Card Market,” *American Economic Review*, 1991, 81 (1), 50–81.
- Bar-Gill, Oren**, “The Law, Economics and Psychology of Subprime Mortgage Contract,” *Cornell Law Review*, 2009, 94, 1073–1151.
- **and Ryan Bubb**, “Credit Card Pricing: The CARD Act and Beyond,” *Cornell Law Review*, 2012, 97, 967–1018.

- Barber, Brad M., Terrance Odean, and Lu Zheng**, “Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows,” *Journal of Business*, 2005, 78 (6), 2095–2119.
- Bergstresser, Daniel, John M. R. Chalmers, and Peter Tufano**, “Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry,” *Review of Financial Studies*, 2009, 22 (10), 4129–4156.
- Berk, Jonathan B. and Jules H. van Binsbergen**, “Measuring Managerial Skill in the Mutual Fund Industry,” 2012. Working Paper.
- and **Richard C. Green**, “Mutual Fund Flows and Performance in Rational Markets,” *Journal of Political Economy*, 2004, 112 (6), 1269–1295.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian**, “How Does Simplified Disclosure Affect Individuals’ Mutual Fund Choices?,” in David A. Wise, ed., *Explorations in the Economics of Aging*, Chicago: University of Chicago Press, 2011, pp. 75–96.
- Besley, Timothy and Andrea Prat**, “Handcuffs for the Grabbing Hand? Media Capture and Government Accountability,” *American Economic Review*, 2006, 96 (3), 720–736.
- Bhattacharya, Utpal, Andreas Hackethal, Simon Kaesler, Benjamin Loos, and Steffen Meyer**, “Is Unbiased Financial Advice to Retail Investors Sufficient? Answers from a Large Field Study,” *Review of Financial Studies*, 2012, 25 (4), 975–1032.
- Bolton, Patrick, Xavier Freixas, and Joel Shapiro**, “The Credit Ratings Game,” *Journal of Finance*, 2012, 67 (1), 85–111.
- Borenstein, Severin, Jeffrey K. Mackie-Mason, and Janet S. Netz**, “Exercising Market Power in Proprietary Aftermarkets,” *Journal of Economics & Management Strategy*, 2000, 9 (3), 157–188.
- Bucks, Brian and Karen Pence**, “Do Borrowers Know Their Mortgage Terms?,” *Journal of Urban Economics*, 2008, 64 (2), 218–233.
- Cain, Daylian M., George Loewenstein, and Don A. Moore**, “The Dirt on Coming Clean: Perverse Effects of Disclosing Conflicts of Interest,” *Journal of Legal Studies*, 2005, 34 (1), 1–25.
- Carhart, Mark M.**, “On Persistence in Mutual Fund Performance,” *Journal of Finance*, 1997, 52 (1), 57–82.
- Chalmers, John and Jonathan Reuter**, “What is the Impact of Financial Advisors on Retirement Portfolio Choices and Outcomes?,” 2012. Working Paper.

- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, 99 (4), 1145–1177.
- Cho, In-Koo and David Kreps**, “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, May 1987, 102 (2), 179–222.
- Choi, James J., David Laibson, and Brigitte C. Madrian**, “\$100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans,” *Review of Economics and Statistics*, 2011, 93 (3), 748–763.
- Christoffersen, Susan E. K., Richard Evans, and David K. Musto**, “What Do Consumers’ Fund Flows Maximize? Evidence from Their Brokers’ Incentives,” *Journal of Finance*, 2013, 68 (1), 201–235.
- Congdon, William J., Jeffrey R. Kling, and Sendhil Mullainathan**, *Policy and Choice: Public Finance Through the Lens of Behavioral Economics*, Brookings Institution Press, 2011.
- Cruickshank, Don**, “Competition in UK Banking: A Report to the Chancellor of the Exchequer,” Technical Report, HM Treasury 2000.
- Dahremöller, Carsten**, “Unshrouding for Competitive Advantage,” *Journal of Economics & Management Strategy*, 2013, 22 (3), 551–568.
- Daniel, Kent, David Hirshleifer, and Siew Hong Teoh**, “Investor Psychology in Capital Markets: Evidence and Policy Implications,” *Journal of Monetary Economics*, 2002, 49, 139–209.
- Del Guercio, Diane and Jonathan Reuter**, “Mutual Fund Performance and the Incentive to Generate Alpha,” *Journal of Finance*, 2012. Forthcoming.
- DellaVigna, Stefano and Ulrike Malmendier**, “Contract Design and Self-Control: Theory and Evidence,” *Quarterly Journal of Economics*, 2004, 119 (2), 353–402.
- Diamond, Peter A.**, “A Model of Price Adjustment,” *Journal of Economic Theory*, 1971, 3, 156–168.
- Dranove, David and Ginger Zhe Jin**, “Quality Disclosure and Certification: Theory and Practice,” *Journal of Economic Literature*, 2010, 48 (4), 935–963.
- Duarte, Fabian and Justine S. Hastings**, “Fettered Consumers and Sophisticated Firms: Evidence from Mexico’s Privatized Social Security Market,” 2012. Working Paper.
- Eliaz, Kfir and Ran Spiegler**, “Contracting with Diversely Naive Agents,” *Review of Economic Studies*, 2006, 73 (3), 689–714.
- and —, “Consumer Optimism and Price Discrimination,” *Theoretical Economics*, 2008, 3, 459–497.

- Ellison, Glenn**, “A Model of Add-On Pricing,” *Quarterly Journal of Economics*, 2005, *120* (2), 585–637.
- and **Sara Fisher Ellison**, “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica*, 2009, *77* (2), 427–452.
- Evans, David S. and Richard Schmalensee**, *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, second ed., Cambridge and London: MIT Press, 2005.
- Eyster, Erik and Matthew Rabin**, “Cursed Equilibrium,” *Econometrica*, 2005, *73* (5), 1623–1672.
- , — , and **Dimitri Vayanos**, “Financial Markets where Traders Neglect the Informational Content of Prices,” 2013. Working Paper.
- Farrell, Joseph and Paul Klemperer**, “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” in “Handbook of Industrial Organization,” Vol. 3 2007, chapter 31, pp. 1967–2072.
- 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.
- , — , **Deyuan Li, Hongyi Li, Sidney Resnick, and Casper G. de Vries**, “The Impact of Competition on Prices with Numerous Firms,” 2012. Working Paper.
- Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny**, “Money Doctors,” 2012. Working Paper.
- Gerardi, Kris, Lorenz Goette, and Stephan Meier**, “Financial Literacy and Mortgage Outcomes,” 2009. Paper presented at the ASSA meetings.
- Gil-Bazo, Javier and Pablo Ruiz-Verdú**, “The Relation between Price and Performance in the Mutual Fund Industry,” *Journal of Finance*, 2009, *64* (5), 2153–2183.
- Gorton, Gary and Andrew Winton**, “Financial Intermediation,” in George Constantinides, Milt Harris, and Rene Stulz, eds., *Handbook of the Economics of Finance*, Amsterdam: North Holland, 2003, pp. 431–552.
- Gottlieb, Daniel and Kent Smetters**, “Narrow Framing and Life Insurance,” 2012. Working Paper.
- Grubb, Michael D.**, “Selling to Overconfident Consumers,” *American Economic Review*, 2009, *99* (5), 1770–1807.

- , “Consumer Inattention and Bill-Shock Regulation,” 2012. Working Paper.
- Gruber, Martin J.**, “Another Puzzle: The Growth in Actively Managed Mutual Funds,” *Journal of Finance*, 1996, 51 (3), 783–810.
- Gurun, Umit G., Gregor Matvos, and Amit Seru**, “Advertising Expensive Mortgages,” 2013. Working Paper.
- Hackethal, Andreas, Michael Haliassos, and Tullio Jappelli**, “Financial Advisors: A Case of Babysitters?,” *Journal of Banking & Finance*, 2012, 36 (2), 509–524.
- , **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 Köszegi**, “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,” in Dan Sjöblom, ed., *The Pros and Cons of Consumer Protection*, Stockholm: Swedish Competition Authority, 2012, pp. 44–76.
- , — , and — , “Inferior Products and Profitable Deception,” 2012a. Working Paper.
- , — , and — , “Exploitative Innovation,” 2012b. Working Paper.
- Inderst, Roman and Marco Ottaviani**, “Misselling through Agents,” *American Economic Review*, 2009, 99 (3), 883–908.
- and — , “Competition through Commissions and Kickbacks,” *American Economic Review*, 2012, 102 (2), 780–809.
- and — , “Financial Advice,” *Journal of Economic Literature*, 2012, 50 (2), 494–512.
- and — , “How (Not) to Pay for Advice: A Framework for Consumer Financial Protection,” *Journal of Financial Economics*, 2012, 105, 393–411.
- and — , “Sales Talk, Cancellation Terms, and the Role of Consumer Protection,” *Review of Economic Studies*, 2013, 80 (3), 1002–1026.
- Ko, Jeremy K. and Jared Williams**, “Disclosure and Price Regulation in a Market with Shroudable Costs,” 2011. Working Paper.
- Kosfeld, M and U Schüwer**, “Add-on Pricing, Naive Consumers, and the Hidden Welfare Costs of Education,” 2011. Working Paper.

- 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.
- Lacko, James M. and Janis K. Pappalardo**, “Improving Consumer Mortgage Disclosures,” *Federal Trade Commission Staff Report*, 2007.
- Laibson, David and Leeat Yariv**, “Safety in Markets: An Impossibility Theorem for Dutch Books,” 2007. Working Paper.
- Laibson, David I., Andrea Repetto, and Jeremy Tobacman**, “Estimating Discount Functions with Consumption Choices over the Lifecycle,” 2007. NBER Working Paper #13314.
- Lal, Rajiv and Carmen Matutes**, “Retail pricing and advertising strategies,” *Journal of Business*, 1994, 67 (3), 345–370.
- Li, Sanxi, Martin Peitz, and Xiaojian Zhao**, “Vertically Differentiated Duopoly with Unaware Consumers,” *Mathematical Social Sciences*, 2012. Forthcoming.
- Lizzeri, Alessandro**, “Information Revelation and Certification Intermediaries,” *RAND Journal of Economics*, 1999, 30 (2), 214–231.
- Malkiel, Burton G.**, “Returns from Investing in Equity Mutual Funds 1971 to 1991,” *Journal of Finance*, 1995, 50 (2), 549–572.
- Malmendier, Ulrike and Devin Shanthikumar**, “Are Small Investors Naive about Incentives?,” *Journal of Financial Economics*, 2007, 85 (2), 457–489.
- Meier, Stephan and Charles Sprenger**, “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal: Applied Economics*, 2010, 2 (1), 193–210.
- 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,” 2010. Working Paper.
- Murooka, Takeshi**, “Deception under Competitive Intermediation,” 2013. Working Paper.
- O’Donoghue, Ted and Matthew Rabin**, “Doing It Now or Later,” *American Economic Review*, 1999, 89 (1), 103–124.
- Piccione, Michele and Ran Spiegler**, “Price Competition under Limited Comparability,” *Quarterly Journal of Economics*, 2012, 127 (1), 97–135.

- Reinganum, Jennifer F.**, “The Timing of Innovation: Research, Development, and Diffusion,” in Richard Schmalensee and Robert D. Willig, eds., *Handbook of Industrial Organization*, Vol. 1, Amsterdam; Oxford and Tokyo: North-Holland, 1989, pp. 849–908.
- Shapiro, Carl**, “Aftermarkets and Consumer Welfare: Making Sense of Kodak,” *Antitrust Law Journal*, 1995, 63 (2), 483–511.
- Shilony, Yuval**, “Mixed Pricing in Oligopoly,” *Journal of Economic Theory*, 1977, 14, 373–388.
- Skiba, Paige Marta and Jeremy Tobacman**, “Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing, Repayment, and Default,” 2008. Vanderbilt Law and Economics Research Paper No. 08-33.
- Spiegler, Ran**, “Competition over Agents with Boundedly Rational Expectations,” *Theoretical Economics*, 2006, 1, 207–231.
- , “The Market for Quacks,” *Review of Economic Studies*, 2006, 73 (4), 1113–1131.
- , *Bounded Rationality and Industrial Organization*, Oxford University Press, 2011.
- Stahl, Dale O.**, “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review*, 1989, 79 (4), 700–712.
- Stoughton, Neal, Youchang Wu, and Josef Zechner**, “Intermediated Investment Management,” *Journal of Finance*, 2011, 66 (3), 947–980.
- Varian, Hal R.**, “A Model of Sales,” *American Economic Review*, 1980, 70 (4), 651–659.
- Warren, Patrick L. and Daniel H. Wood**, “Will Governments Fix What Markets Cannot? The Positive Political Economy of Regulation in Markets with Overconfident Consumers,” 2010. Working Paper.
- Woodward, Susan E. and Robert E. Hall**, “Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence,” *American Economic Review*, 2012, 102 (7), 3249–3276.