CHAPTER 6

Behavioral Industrial Organization*

Paul Heidhues[*,1], **Botond Kőszegi**[†]

[*]DICE, Heinrich-Heine Universität, Düsseldorf, Germany
[†]Central European University, Budapest, Hungary
[1]Corresponding author. E-mail address: heidhues@dice.uni-duesseldorf.de

Contents

* We thank Mats Köster, Frank Schlütter, and especially Stefano DellaVigna for helpful comments.

Handbook of Behavioral Economics, Volume 1
ISSN 2352-2399, https://doi.org/10.1016/bs.hesbe.2018.07.006

1. INTRODUCTION

In industrial organization's influential long history, researchers have devoted tremendous attention to thinking about the precise details of the economic environment firms and regulators are operating in, and how these details influence firm behavior, market outcomes, and possibilities for welfare-increasing intervention. Thousands of papers point out ways in which the complementarity versus substitutability of products, the information, technological capabilities, and commitment power of firms and regulators, and the timing of market interactions and interventions crucially affect consumer and producer welfare. Yet until recently, researchers have devoted a comparatively tiny amount of attention to thinking about the individuals driving market outcomes: how consumers and managers perceive the environment they are operating in, what personal goals and expectations they have in going to the market, and how these goals interact with their individual mental capabilities to shape behavior. In this chapter, we review research that aims to make progress on the latter questions, broadly summarized under the umbrella Behavioral Industrial Organization. Only ten years ago, Editors of the third volume of the Handbook of Industrial Organization did not consider it essential to cover the topic, and indeed at that stage research on it was only in its infancy. Since then, there has been an explosion of research, warranting a comprehensive review.

As it is the case with behavioral economics in general, defining what we view as behavioral industrial organization—and therefore what we include in this chapter—entails some difficult and often debatable subjective judgments. We have converged on two criteria that the research we review must satisfy. First, consistent with what is usually considered industrial organization, the work should explore economically important market interactions between firms, or between firms and consumers. This means that we exclude topics such as contracting inside a firm, auctions and mechanism design, public finance that does not involve non-trivial models of firm behavior, or experimental work where market interactions are simulated in the lab. Second, consistent with what we (but not all researchers) consider psychology and economics, we require that the central feature of the decisionmaking model at the heart of a contribution be psychologically well-founded. This excludes, for instance, applications of rational inattention where unfounded features of the attention-cost function are central for the results, or macroeconomics-style models of menu costs whose source is unclear. We also do not review issues that are too context-specific to draw general conclusions from.

Even with the above relatively strict inclusion criteria, the research we review is already large and quickly expanding. This development strikes us as quite natural, given that (as we will argue) the insights are relevant for understanding outcomes and welfare in some of the most important economic markets, as well as evaluating possible policies that pertain to those markets. Consumers' systematic failure to fully understand offers, or their mistakes in predicting their own behavior, appear helpful in explaining why credit cards have high interest rates, why bank accounts charge high overdraft

fees, or why life-insurance contracts have front-loaded fees—and suggest that we should think about (but not automatically jump to) regulating such charges. Consumers' limited attention, combined with firms' limited incentive to educate, but often substantial incentive to obfuscate, appear useful for understanding why confusion still reigns free in many consumer financial markets, such as those for mortgages and mutual funds—and why strict disclosure regulations have not solved the problem. And loss aversion appears helpful in explaining patterns in pricing under oligopolistic competition in a variety of consumer retail markets, such as why sellers charge the same price for differentiated products—questioning the received wisdom that such patterns reflect collusion by firms.

Importantly, although we require models to be well-founded, we do not discuss evidence for individual-decisionmaking models; at most, we mention such evidence as motivation and refer readers to other chapters in the Handbook. We include evidence only if it is specifically about market interaction. And although we require that some evidence or economic argument support the decisionmaking model used, we do not take the view that a researcher must conclusively rule out the existence of a classical account for a market before considering alternatives. Model uncertainty, or current unavailability of data to distinguish models, does not justify sticking with the classical model. Furthermore, developing and then testing distinct market implications of well-founded behavioral-economics models is a fruitful indirect way of going about testing these models. As a simple example, observing firms' strict preference to hide some price components is inconsistent with models of rational consumers—where adverse selection dictates that consumers would anticipate hidden prices to be high—and hence provides indirect evidence for consumer naivete. We will point out when the theoretical results we discuss have such implications.

In the rest of this introduction, we outline the specific topics we will cover. Our review is organized around theoretical principles rather than empirical findings or methods. The reason is simple: at the current stage of the literature, the theoretical side of behavioral industrial organization is far more developed than the empirical side. We discuss empirical work where it fits within the conceptual organization. We hope that empirical researchers will soon come to study more industrial-organization issues. Indeed, that pattern of research development, whereby a spate of theoretical research is followed by an emphasis on empirical testing, would be consistent with what happened in classical industrial organization.

A very substantial part of behavioral industrial organization studies situations in which consumers make systematic mistakes, or are "naive," in assessing the value they will derive from a product. Usually, this is formalized by assuming that consumers ignore part of what they will pay for the product. Because this simple possibility of "hidden prices" raises a host of different issues, we devote three sections to it. We begin in Section 2 with what can be viewed as the basic market implications of hidden prices:

we study situations in which a firm takes advantage of consumers' mistakes to make its own product look more attractive, but in doing so does not affect how consumers value other firms' products. Under the simplest of circumstances, competition protects consumers from the effects of their own naivete, and in fact competition does more for naive consumers than for rational consumers. But in a host of more complicated situations naivete has distributional as well as efficiency implications. Competition often leads to a redistribution from naive to sophisticated consumers, and leads to too many consumers participating in the market. And independently of the degree of competition, consumer naivete induces firms to sell products or contracts with inefficient features geared toward exploiting naivete. We also show how mistakes that have been documented in other settings can lead to hidden prices in a market setting, and how to detect consumer mistakes from market data.

Analogously to classical preference-based price discrimination—in which firms treat consumers with different preferences differently—in Section 3 we review work on naivete-based discrimination—in which firms treat consumers with different degrees of naivete differently. Naivete-based discrimination raises theoretically novel issues. With rational consumers, a consumer's (potentially stochastic) behavior and her ex-ante beliefs about her behavior coincide. With naive consumers, they may not, raising two new motives for discrimination. First, a firm may want to discriminate between consumers who have the same ex-post preferences (and hence behavior) but different ex-ante beliefs. Since consumers with different beliefs choose from available offers in a different way, it is possible to induce self-selection among them. This leads to second-degree price discrimination. Second, a firm may want to discriminate between consumers who have the same ex-ante beliefs but behave differently ex post. Since consumers with the same ex-ante beliefs (and ex-ante preferences) always choose from available options in the same way, it is impossible to induce self-selection among them. Hence, discrimination must be based on other information. This leads to third-degree price discrimination.

In Section 4, we consider situations in which a firm can affect not only a consumer's valuation of its own product, but also her valuation of competing products. Such perception externalities can arise because the firm can educate or obfuscate the products on offer, or manipulate what the consumer considers. The research we review establishes that firms often have a limited incentive to educate consumers about hidden fees, and—especially when competition is fierce—they may have a strong incentive to obfuscate hidden fees. Perception externalities also arise if the set of products on offer affects how consumers trade off different features of the products. Finally, perception externalities often arise if consumers have limited attention, so that one firm's product can draw away costly attention from other offers.

In Section 5, we turn to considering markets where both firms and consumers are rational, but consumers have different preferences from those typically assumed in industrial organization. We focus on three well-studied aspects of consumer preferences.

When consumers are loss averse, they are very (first-order) averse to bearing risk. As a result, firms often have an incentive to shield consumers from economic risk, resulting in reduced price variation, flat fees, and fewer choices. Nevertheless, loss aversion can also induce firms to introduce risk into an otherwise riskless environment, resulting in temporary price discounts or limited-availability sales. When consumers have present bias or temptation disutility, they prefer to commit their future behavior, and profit-maximizing firms have an incentive to fill this demand. But the demand for commitment may be limited due to uncertainty, and market conditions may place constraints on commitment possibilities, resulting in a rich set of tradeoffs. And when consumers purchase products not just for consumption benefits, but also to show off their wealth or good taste, market competition looks very different than when firms supply classical products.

All of the above topics pertain to situations in which firms are identical to those typical in classical industrial organization: they are rational and profit-maximizing. In Section 6, we discuss the literature on how psychological phenomena may affect firms. Just like consumers, managers can be subject to mistakes or limited attention, and therefore they do not always make optimal decisions. This manifests itself in mistakes in pricing and investment decisions and in gauging competitors' behavior, which may occur for both small and large firms. Psychological considerations are also important for understanding contracting between firms and therefore the theory of the firm, as a contract can induce a change in preferences that affects how parties interact at a later stage. And firms may have—or consumers may expect them to have—motivations beyond profit maximization, especially regarding other market measures or social causes.

In Section 7, we discuss research addressing possible policy implications of psychological phenomena. As our overview in Sections 2 through 4 indicates, it is in general difficult to leverage market forces to eliminate the effects of consumer mistakes. Because heavy interventions can create large welfare losses if used in the wrong situations and because many researchers and policymakers dislike heavy interventions on principle, a lot of the emphasis among applied researchers has been on soft interventions—interventions that help naive consumers without changing their choice set or without hurting sophisticated consumers. We point out, however, that once equilibrium considerations are taken into account, soft interventions are unlikely to be available. Furthermore, the most obvious approach for dealing with naive consumers, education or improved disclosure, may not work or may have undesirable side-effects. We therefore consider heavier interventions that regulate the types of products or contracts firms can sell. Furthermore, we review research on how considerations in behavioral industrial organization affect classical policy recommendations. Most importantly, researchers studying puzzling market phenomena from a classical industrial-organization perspective have sometimes attributed those phenomena to welfare-decreasing firm behavior such as predation or collusion. New work has pointed out that plausible explanations based on psychologi-

cally motivated consumers may also be available, so that a regulatory response may not be justified.

2. BASIC ECONOMICS OF HIDDEN PRICES

By far the largest topic in behavioral industrial organization concerns situations in which rational profit-maximizing firms who fully understand the environment they are operating in—including being able to correctly forecast (average) consumer behavior—interact with consumers prone to making systematic mistakes. Both because of the volume of research and for conceptual reasons, we break up the literature into three parts. In this section, we consider situations in which a firm competes solely by trying to take advantage of consumer mistakes to make its own product look good, taking the consumer's perceived outside option as given. In the next section, we review research on how firms can price discriminate between heterogenous consumers when some consumers make mistakes. Finally, we consider various ways in which firms can manipulate the mistakes consumers are making, so that they can affect the perceived valuation of alternative products.

2.1 Framework and safety-in-markets benchmark

We introduce a bare-bones reduced-form model that will allow us to discuss, in a single framework, insights from a wide range of papers employing different approaches. We start with the simplest benchmark model, and in the subsequent subsections add features to introduce themes that have been explored in the literature.

The additional price. Generalizing a framework introduced by Gabaix and Laibson (2006), the key assumption of our model is that a firm can leverage consumer mistakes to charge an "additional price" a that naive consumers ignore when making purchase decisions, but that they pay if they purchase. That many consumers pay unexpected charges is supported by a large body of evidence cited in the chapter Behavioral Household Finance of this handbook. The literature discusses three basic types of systematic mistakes that can lead consumers to pay unexpected charges. First and most fundamentally, consumers may misunderstand the contract itself and, hence, misestimate the amount they will end up paying when accepting a product or contract.[1] Second, in

[1] A common reaction we have encountered in seminars and referee reports in response to this assumption is that (at least in developed economies) prices are disclosed, so it is unreasonable to assume that consumers systematically underestimate the price. In our view, such a reaction is misplaced: whether a disclosed price is an understood price is an empirical question. And indeed, the evidence suggests that it is often not. For instance, different types of empirical evidence suggest that—despite disclosure regulations—investors do not fully understand the management fees of mutual funds, and that they appreciate front loads better. See the chapter Behavioral Household Finance of this handbook.

what Bar-Gill and Ferrari (2010) refer to as use-pattern mistakes, consumers may incorrectly forecast their own future behavior and hence the cost associated with accepting a product or contract. Third, consumers may misunderstand some statistical features of the world.

An alternative possibility is that consumers misunderstand not the price, but the quality or value of the product. For most of the issues we discuss, this generates similar effects as a price misperception: the consumer overestimates the net value she gets from purchasing, and the firm makes extra profits. The only difference is that in the case of value misperception, the extra profits come from lower costs or selling more rather than higher revenues.

Some authors posit exogenously that firms can charge an additional price. Other work, which we review in more detail below, derives the additional price from the interaction of profit-maximizing firms with consumers who make "primitive" behavioral mistakes documented in other settings. Taken together, researchers have identified countless foundations for an additional price. We briefly mention three very different examples. (1) Armstrong and Vickers (2012) argue that some consumers simply ignore overdraft charges, choosing bank accounts as if the overdraft charges did not exist. (2) Grubb (2009) assumes that mobile-phone consumers underestimate the variance of their demand for calling. Providers respond by making the price convex in minutes of calling. With such plans, consumers' mistake leads them to underestimate the expected price. (3) In Spiegler (2006b), "quacks" who produce no value offer treatments to customers whose outcome is random. In a law-of-small-numbers type of statistical mistake, a customer believes a quack offers successful treatments whenever the treatment she observes was successful, and otherwise believes the quack's treatment to lead to certain failure. The observation of a successful treatment, hence, leads consumers to be willing to pay for a useless service. As will be clear from our analysis, for many purposes the source of the additional price is unimportant, and it is very convenient to work with a reduced-form model. But to understand some issues, including exploitation distortions and screening discussed below, it is necessary to model the underlying consumer behavior. Any researcher aiming to make progress regarding the economics of hidden pricing must carefully consider her research agenda to decide whether a reduced-form or more well-founded model is called for.

All existing models of consumer mistakes posit that in addition to misunderstanding the price or the product, consumers also make a strategic mistake. In particular, consumers do not make inferences about their own naivete from the products and contracts offered by firms. While there is little evidence for this assumption and investigating relaxations would be useful—for instance, consumers may become suspicious of overly attractive-looking deals—we feel that it is highly plausible as a starting point. Consumers who do not understand basic features of the product they are getting are even less likely to make sophisticated equilibrium inferences from firm behavior.

Market framework. Following Heidhues and Kőszegi (2017), our bare-bones model adds the above possibility of charging additional prices into a duopoly pricing model with horizontally-differentiated products. The duopolists are located at the endpoints $l = 0$ and $l = 1$ of the unit interval, and offer their products to consumers who are interested in buying at most one unit. Consumers are uniformly distributed over the unit interval, and a consumer located at y incurs a disutility or "transportation cost" of $t|y - l|$ when buying product l, where $t > 0$ is a product-differentiation parameter that determines a firm's market power. A consumer has gross utility of v of acquiring the product, and her outside option has gross utility 0; the outside option, however, is available only at the endpoints of $[0, 1]$, so that for consumer y it has utility $-t \min\{y, 1 - y\}$.[2] Both firms have identical marginal costs c and simultaneously set anticipated prices $f_l \in \mathbb{R}$ and additional prices $a_l \in [0, a_{max}]$, where $a_{max} > 0$. Crucially, despite the fact that consumers end up paying a_l when purchasing the product, they ignore this price component when making purchase decisions.

We suppose firms have a correct understanding of consumers' behavior, and analyze *symmetric pure-strategy Nash equilibria* of the firms' pricing game. Depending on firms' market power t, a firm that raises its anticipated price f_l slightly above the equilibrium level either loses consumers to its rival or the outside option. We refer to the former case as an imperfectly competitive market and the latter case as a monopolistic market. To understand the effects of consumer mistakes, it is useful to compare the equilibrium outcome of a model with naive consumers to that of a model in which all consumers are "rational" in that they fully understand that they end up paying $f_l + a_l$ when accepting the contract of firm l.

High additional price—but safety in markets. Because consumers ignore the additional price when shopping, both firms set $a_l = a_{max}$, so that the profits on the additional price can be very high. Indeed, numerous authors in the literature propose that the high prices we observe in some markets are due to consumers not paying (full) attention to those prices. In early contributions, for instance, Ausubel (1991) and DellaVigna and Malmendier (2004) contend that the high interest rates for credit cards are due to the presence of borrowers who underestimate borrowing and hence underappreciate the importance of interest rates. And Bar-Gill (2009) argues that the escalating payments in subprime mortgage contracts are consistent with borrowers underestimating future costs.

Despite the high additional prices, however, in the current model an important "safety-in-markets" benchmark result obtains: Suppose the good is socially valuable

[2] This specification of the outside option follows Benabou and Tirole (2016). In contrast to the standard formulation (Hotelling, 1929), the product-differentiation parameter t only impacts the level of competition and not the attractiveness of purchasing relative to the outside option. It, thus, allows us to single out the effect of competition on outcomes.

($v > c$) and the market imperfectly competitive. Then, consumers' equilibrium welfare is unaffected by consumer naivete. Based on the same logic as in switching-cost (e.g., Farrell and Klemperer, 2007) and loss-leader (e.g., Lal and Matutes, 1994) models, competition leads firms to hand the profits from the additional price back to consumers by lowering the anticipated price. Hence, consumers end up paying the same amount for the service as they would if they fully understood the contract. Although in a completely different framework, the idea that competition can protect irrational consumers already appears in Laibson and Yariv (2007), and indeed variants of the above benchmark result are derived in a number of papers (e.g., Grubb, 2015a).

The safety-in-markets argument depends crucially on the market being (at least imperfectly) competitive. In the monopolistic case, a firm in equilibrium makes consumers indifferent between accepting and rejecting its offer. Unforeseen additional charges do not affect consumers' willingness to accept an offer, and hence consumer naivete—by allowing a monopolistic firm to raise its total price $f_l + a_l$ and still sell its product—hurts consumer welfare. Consumer naivete, therefore, strengthens the case for competition policy as a means of increasing consumer welfare: moving from monopoly to perfect competition ($t = 0$) raises consumer welfare from 0 to $v - c$ in the rational case, and from $-a_{max}$ to $v - c$ in the naive case.

2.2 Limitations on competition for naive consumers

The safety-in-markets result relies on competition returning *all* profits from the additional price to consumers. We review a number of economically important reasons that—even if the market is competitive—this may not be the case. In later subsections, we discuss further limitations that arise when consumers are heterogenous in their understanding of offers.

Underweighting of all price components. One important feature of our simple model and others that emphasize safety in markets is that there is some element of the contract—here the anticipated price—that allows for transferable utility and that consumers fully take into account. If a consumer instead underestimates all elements of the total price by a given percentage—due to underestimating the probability of purchase, as in Johnson's (2017) model of unplanned purchases, or ignoring a tax to be added to the price, as in Chetty et al.'s (2009) empirical findings—then she puts too much weight in her purchase decision on how well the product matches her taste, and values price reductions less than a rational consumer. Even in an imperfectly competitive environment, this is equivalent to an increase in transportation costs, and hence leads to higher prices.

Price floors. Competition can also fail to return profits from the additional price in full because there are limits to the extent to which firms can cut the anticipated

price. Some authors simply posit that negative prices are infeasible (e.g., Armstrong and Vickers, 2012; Ko, 2012), but there are also a number of models in which price floors arise endogenously. In Miao's (2010) duopoly model, firms sell a basic good—such as a printer—and an add-on service—such as a new cartridge—and a firm can choose to make its printer incompatible with the cartridges of other firms. Naive (or myopic) consumers do not take the need for cartridges into account, so that—by the same logic as above—firms price cartridges at consumers' valuation. But if a firm sets its printer price too low, a consumer would be better off buying a new printer instead of a cartridge. Hence, the monopoly price in the aftermarket effectively creates a price floor in the primary-good market. Similarly, in Michel (2017) consumers at the point-of-sale of some basic product—like an electronic device—decide whether to buy an extended warranty. Naive consumers do not take the warranty terms into account when selecting a store, and underestimate their return costs and, hence, overvalue an extended warranty. Because buying multiple units of the basic good partially ensures against needing a replacement, firms can only earn profits from selling the warranty if the base good is not priced too low, inducing a price floor below which firms cannot cut the price.

In Heidhues et al.'s model of the mutual-fund industry (2017), funds choose front loads investors incur when buying a share in the mutual fund, and which they fully understand. In addition, mutual funds choose management fees investors ignore. Thinking of these management fees as being the additional price, our bare-bones model predicts that new investors should receive a "signing bonus" or negative front load. Payment of such a signing bonus, however, is ruled out through part of the Investment Company Act of 1940, which forbids favoring new investors relative to existing ones, and thereby effectively requires front loads to be non-negative. Similarly, a number of countries require supermarkets to sell their products above the wholesale price, implying that loss leaders—be they due to rational (Lal and Matutes, 1994) or naive (Johnson, 2017) consumers—cannot be sold below their marginal cost even if firms would want to do so in an equilibrium absent price floors.

In the credit-card model of Heidhues et al. (2017), card issuers offer contracts containing an anticipated price as well as an interest rate naive (time-inconsistent) consumers ignore. Consumers receive a convenience benefit from using the card, but due to their naivete pay unanticipated interest. When indifferent between card offers, each consumer decides whether to get a card based on some exogenously given order among cards, and gets multiple cards if doing so strictly increases utility. Similarly, when deciding on which card to charge, the consumer uses the same exogenous preference order whenever she is indifferent. These assumptions imply that any consumer who prefers a firm's card will get the card if the firm charges an anticipated price of zero. As a result, a situation in which credit-card companies earn unanticipated interest, charge an anticipated price of zero, and earn positive profits is an equilibrium: any additional consumers a negative

anticipated price attracts will not use the firm's card and hence these consumers are unprofitable. Because consumers can multi-home, therefore, firms act as if they were facing a price floor of zero.

Finally, consumer suspicion can also give rise to a price floor. Consumers seeing a low price may begin to wonder just how the price could be so low, and come to believe that "there must be a catch," leading them to refrain from buying. This possibility prevents firms from setting overly low anticipated prices. Perversely, it is exactly consumers' distrust that creates a price floor and enables firms to earn profits in equilibrium; if consumers were entirely unsuspecting, firms would compete more fiercely.[3]

Failure in comparing offers. When all firms price at a price floor, consumers make purchase decisions based on non-price considerations—either based on their tastes, or, for homogenous products, completely randomly. The same consideration arises when consumers cannot compare prices. Evidence that consumers often fail to choose the best price, and the implications for firm behavior, are discussed in detail in Grubb (2015b), and we will also summarize such models in Section 4. As will be clear, the failure to choose the best price can also limit competition for consumers who pay an additional price.

Just like consumers' inability to compare prices, consumers' inability to judge product values can also lower competition. In particular, if consumers with homogenous tastes overvalue different homogenous products by different amounts, the resulting perceived product differentiation enables homogenous firms to earn positive profits. Spiegler (2006b) demonstrates this in a model of "quacks." Patients can acquire a treatment from n profit-maximizing healers, with a patient's true recovery rate being the

[3] While this has not been fully modeled and explored, we provide a sketch based on Heidhues et al. (2012b). Consumers—who may naively think that existing regulations are likely prevent firms from collecting an additional price—a priori think that firm l is restricted to charging $a_l = 0$ with probability $1-\epsilon$, and can charge any additional price $a_l \in [0, a_{max}]$ with probability ϵ. Consumers suppose that the probability of a firm being the "deceptive" type is drawn independently across firms.

Customers believe that all firms' marginal costs are $\tilde{c} \in (c - a_{max}, v)$, and that firms are playing a perfect Bayesian equilibrium given those marginal costs. In fact, however, no such strict regulation of the additional price exists, so that firms are playing a simultaneous-move game in which they choose contracts (f_l, a_l) with $a_l \leq a_{max}$. We want to establish that if $\tilde{c} + \epsilon a_{max} \leq \min\{c, v\}$, there is a Nash equilibrium in this simultaneous-move game between the firms in which all firms choose an up-front price of $\tilde{c}$, implying that $\tilde{c}$ acts as a floor on the up-front price. In this equilibrium, firm l offers a contract in which $f_l = \tilde{c}$, $a_l = a_{max}$, and consumers believe that firm l offers the contract $f_l = \tilde{c}$, $a_l = 0$ if restricted to $a_l = 0$ and the contract $f_l = \tilde{c}$, $a_l = a_{max}$ if not restricted to $a_l = 0$. Whenever firm l deviates and charges an anticipated price strictly below $\tilde{c}$, suspicious consumers—convinced that without additional prices the firm would make losses—believe with probability one that firm l charges an additional price $a_l = a_{max}$. (If the firm charges an anticipated price weakly above $\tilde{c}$, we suppose that consumers do not update their prior and believe that the firm must charge $a_l = 0$ with probability $1-\epsilon$.) Given these consumers' beliefs, when undercutting $\tilde{c}$ consumers correctly predict the firms total price $f_l + a_{max}$, and only buy if this price is below their valuation and is below the expected expenditure $\tilde{c} + \epsilon a_{max}$ they incur when buying from a rival.

same for all healers as well as the outside option. Patients, however, rely on a form of "anecdotal reasoning:" each patient independently samples each treatment as well as the outside option once, and incorrectly believes that an option's success rate equals that in her sample. A firm, hence, has monopoly power over consumers for whom only its own sample treatment was successful, and competes with another firm over consumers for whom both firms' sample treatments were successful. The fact that firms have monopoly power over some consumers enables them to earn positive profits. The tension between exploiting these consumers and undercutting rivals to steal those consumers who are willing to switch leads to a mixed-strategy equilibrium as in Varian (1980).

Evidence. In an influential paper consistent with a model of limited competition for naive consumers, Ausubel (1991) uses multiple empirical approaches to argue that credit-card lending is—from a classical point of view—puzzlingly profitable given the industry's competitiveness by conventional measures of concentration. He finds that credit-card issuers have earned profits that far exceed the rate of return elsewhere in the banking industry. Additional data shows that banks can sell credit-card portfolios at a substantial premium, suggesting that firms expect the profitability to persist. Consistent with our prediction that the additional price is high and not tied to marginal cost, Ausubel also documents that interest rates remained high and stable in a period with large changes in the cost of funds. Ausubel argues that his findings are inconsistent with a model featuring reasonable search or switching costs. In a model of switching costs, the profit a firm can make on a consumer ex post is lower than the consumer's switching cost. Yet the profits lenders earn are far higher than any reasonable switching cost. Instead, Ausubel proposes that a specific form of consumer irrationality, whereby borrowers expect not to borrow and then do, may be involved. In a similar vein, Bar-Gill (2004) points out that basic features of the credit-card market are consistent with consumers underestimating borrowing costs. Adding a price floor or other phenomenon that limits competition in the annual fee and perks accounts for the profitability Ausubel finds.

2.3 Heterogenous naivete and distributional effects

The previous subsection demonstrated how the exploitation of naive consumers can affect the distribution of welfare between firms and consumers. We now show that if consumers are heterogenous in their understanding of products, then the exploitation of naive consumers can also affect the distribution of welfare between consumers, and has further effects on the distribution of welfare between firms and consumers.

To isolate the implications of heterogeneity in *naivete* rather than heterogeneity in other things, most models in the literature assume that consumers are homogenous in everything else. Yet consumers could differ in other things, such as their

taste. The implications of such multidimensional heterogeneity are relatively unexplored.[4]

Cross-subsidy from naive to sophisticated consumers. To capture heterogeneity in our bare-bones framework, suppose that in addition to the naive consumers, there are some sophisticated consumers—consumers who fully understand the contracts they are offered—in the population. Each consumer is, independently of her taste, naive with probability α and sophisticated with probability $1-\alpha$. We begin by supposing that sophisticated consumers anticipate the additional price a_l and can costlessly avoid paying it. Hence, both sophisticated and naive consumers' perceived utility when buying product l is $v-f_l-t|\gamma-l|$, but naive consumer's actual utility is $v-f_l-a_l-t|\gamma-l|$. In the symmetric pure-strategy equilibrium of the game between firms, firms' pricing satisfies $f(\alpha)=\min\{c+t-\alpha a_{max}, v\}$ and $a(\alpha)=a_{max}$.

Since the demand of both sophisticated and naive consumers is independent of a, in equilibrium $a(\alpha)=a_{max}$. As before, in the imperfectly competitive case the expected ex-post profits from attracting a consumer, αa_{max}, are handed back to consumers through a lower anticipated price. Hence, the presence of naive consumers benefits sophisticated consumers, and because the expected handout is less than the additional price naive consumers pay, the presence of sophisticated consumers hurts naive consumers. As first emphasized by Gabaix and Laibson (2006), therefore, naive consumers in effect cross-subsidize sophisticated consumers.[5] In as much as naive consumers are poorer, this can give rise to very adverse distributional effects. For example, in the UK banking industry discussed in detail by Armstrong and Vickers (2012), one can think of the—typically zero—account maintenance fee as the anticipated price, and of the—typically high and by many consumers unanticipated—overdraft fee as the additional price. The overdraft fee is incurred by a fraction of consumers, but of these consumers many incur the fee multiple times; and the subpopulation that does has significantly lower income than the rest. This leads the authors to cite a judge who describes the business model of high overdraft fees together with free banking services as a "reverse Robin Hood exercise."

Arbitrageur-induced price floors. Beyond those discussed in Section 2.2, the presence of sophisticated consumers introduces yet another potential source for a floor on the anticipated price (Armstrong and Vickers, 2012; Ko, 2012), which lowers the extent to which sophisticated consumers benefit from the exploitation of naive consumers, and increases the extent to which firms benefit. Following Heidhues et al. (2012a), consider

[4] Grubb (2015a) allows for such a case, and demonstrates in an example that this can lead to one firm specializing in offering an efficient contract while the rival offers a deceptive one.

[5] In the monopolistic case, no such cross-subsidy occurs. Firms then maximally exploit consumer naivete by setting $f_l=v$ and $a_l=a_{max}$.

our bare-bones model with naive consumers from above but suppose that $v < c$ and for simplicity consider the limit case of homogenous products ($t \to 0$). Furthermore, suppose that in addition to the naive consumers, there are many sophisticated consumers who have a valuation of zero for the product; that is, there are many rational arbitrageurs who accept the contract offer if and only if its price is negative. This effectively induces a price floor as long as firms selling to these arbitrageurs at a negative price cannot recoup the losses from doing so.

At the heart of this arbitrageur-type argument lies the idea that lower anticipated prices attract disproportionally less profitable consumers, a possibility first modeled in Ellison's (2005) model of add-on pricing. Such an adverse attraction effect lowers competition and raises profits, and may be extremely relevant in practice. If more sophisticated consumers are not only better at avoiding contract terms designed to exploit them but also better at comparing prices and selecting the cheapest offer, lowering the anticipated price should often disproportionally attract less profitable consumers (Grubb, 2015a).

2.4 Distortions

The implications of consumer naivete we have identified so far are all distributional. We now discuss distortions that can arise from firms taking advantage of naive consumers. Some of the effects have a parallel in the industrial-organization literature on aftermarket monopolization. This literature assumes that after buying a primary good, consumers are locked into a complementary-good market (aftermarket), and when purchasing the primary good consumers do not observe the (future) price of the complementary good. In such a setting, Shapiro (1995), Hall (1997), and Borenstein et al. (2000) demonstrate that the overly low price in the primary market and the overly high price in the aftermarket lead to inefficiencies. The former parallel our participation distortions, and the latter parallel a special case of our exploitation distortions.

Participation distortions. The fact that naive consumers underestimate the total price of the product (or overvalue its benefit) can have a distortionary effect by inducing consumers to buy even when their value from the product is lower than the production cost. To see this in a trivial example, suppose all consumers are naive in our framework above ($\alpha = 1$) and consider the case of a perfectly competitive market (i.e., the limit as $t \to 0$). Then, if $c - a_{max} < v < c$ consumers buy a wasteful product whose social value is negative.

More generally, in a competitive market with a downward-sloping demand curve a participation distortion always arises for marginal consumers, as the underestimation of the total price affects demand in the same way as a subsidy. In some markets, such as the one for bank accounts considered in Armstrong and Vickers (2012), it is likely to be efficient for most consumers to have an account, so the concern about participation dis-

tortions pales in comparison to other concerns, such as the adverse distributional effects we have discussed in Section 2.3. But for other settings, participation distortions may be very important. For example, in a provocative piece (Heidhues and Kőszegi, 2015), we demonstrate through a simple calibration exercise that the participation distortion in the US credit-card market *may* be enormous—as high as half of the size of the market. As carefully discussed in Grubb (2015c), however, the participation distortion induced by consumer naivete depends on the demand and supply elasticities, and hence requires more research to evaluate.

Brown et al. (2012) identify a mistake in consumer responses to movies that critics are not permitted to review beforehand. The authors document that these "cold openings" are correlated with a pattern of fan disappointment, suggesting that moviegoers do not properly account for the fact that these movies are not reviewed. At the same time, these movies earn more domestic box-office revenues than reviewed movies of similar quality, which—although controlling for endogeneity and selection is difficult—suggests that firms increase demand by profitably exploiting consumers' failure to reason strategically.[6] As consumers treat movies with the same quality differently, again the consumers' misevaluation leads to a participation distortion.

Of course, the participation distortion interacts with the level of competition. If firms have market power and hence charge prices above marginal costs—so that in a classical setting too few consumers would purchase—then the increased participation due to consumers' underestimation of prices may be beneficial. As a specific example, de Meza and Reyniers (2012) show that additional prices can decrease total prices and increase consumer and total welfare in a Cournot model with constant elasticity of demand.[7] Nevertheless, one should not jump to the conclusion that we should allow hidden prices in oligopolistic markets. Indeed, this argument is akin to the suggestion that polluting firms should be allowed to collude to raise price and lower dirty production, or that firms with market power should receive a subsidy to overcome the

[6] Similarly, Mathios (2000) suggests that consumers underinfer low quality from receiving no news regarding the fat content of salad dressing. He looks at the introduction of the Nutrition Labeling and Education Act, which required producers to disclose the fat content. In a classical model with rational consumers and cheap disclosure, even absent regulation all but the highest-fat-content salad dressings should disclose their fat content (Grossman, 1981; Milgrom, 1981). While Mathios finds significant labeling by "low-fat" producers prior to the Act, there is also considerable variation in the fat content of products that are not labeled. After the introduction of the compulsory labeling law, consumers purchase less of the highest-fat-content dressing, indicating that the lack of labeling had an effect on their purchase behavior.

[7] To see heuristically that the total price can decrease, suppose the reduction in the up-front price would exactly offset the additional price. Then each firm's value of a marginal consumer would be the same, but it would have more inframarginal consumers. If—as would be true in the case of linear demand—the derivative of inverse demand would stay the same, the firm would have an incentive to increase its price, suggesting that the symmetric equilibrium price will be higher. But with a constant elasticity demand curve the derivative of inverse demand falls in absolute value, making quantity increases more attractive, potentially leading to lower equilibrium prices.

allocative distortion from overly high prices. We typically think of these as bad ideas, among other things, because they tend to induce excess entry into the undesirable practice.

Indirect distortions from profitable hidden prices. In the presence of a floor on the anticipated price, the potential profits from exploiting naive consumers can induce firms to engage in a number of inefficient behaviors. At a basic level, the quest for these profits can generate overentry into the industry, or lead to excessive marketing efforts such as advertising, mailing of contract offers, or paying commissions to intermediaries.

More subtly, Heidhues et al. (2016) ask a basic question regarding products with hidden prices: where do the hidden prices come from? Coming up with novel ways to exploit consumers, similar to inventing product improvements, presumably is costly and therefore begs the question what a firm's incentive is to engage in such"exploitative innovation." Given that many of these novel exploitative features—particularly in financial markets—are in contract terms that can be easily copied, the incentives to develop these are unclear from a classical perspective. To study this question, Heidhues et al. append a pricing game similar to that in our bare-bones model with a preceding stage in which one firm can make innovations. They focus on the incentives for exploitative innovation—i.e., increasing the maximum additional price the firm can charge—but contrast these with the often-studied incentives for making product improvements that consumers value—i.e., increasing consumers' value for the firm's product. Without a price floor, the incentive for either kind of innovation equals the "appropriable part" of the innovation—the part competitors cannot copy—resulting in similar incentives for exploitative and value-increasing innovations. But in the presence of a binding price floor, the innovation incentives are typically stronger for exploitative than for value-increasing innovations. As a good approximation of contract-clause innovations, consider innovations that any competitor can copy in full. Because in a deceptive market with a binding price floor a higher additional price increases the markup, a firm is willing to engage in such a non-appropriable exploitative innovation. In contrast, because non-appropriable value-increasing innovation does not change the equilibrium prices, and gives no competitive advantage in attracting consumers, a firm has no incentive to engage in it.

Exploitation distortions. When a firm sells to naive consumers, it has an incentive to tilt its business model from providing efficient service to generating revenue from hidden fees. When a firm does so, it offers a contract that—conditional on being accepted—is inefficient. We refer to the resulting distortion as an exploitation distortion. For example, if naive consumers underestimate the importance of roaming fees when signing a contract, mobile-plan providers may charge high roaming fees. Faced

with these high roaming fees when abroad, consumers have an incentive to reduce their amount of calling, generating an inefficiency. And (as we will formalize below) if naive time-inconsistent borrowers pay unanticipated interest on their debt, lenders have an incentive to ramp up this debt in order to collect more unanticipated interest payments, leading them to induce inefficient overborrowing.

Following Heidhues and Kőszegi (2017), we capture exploitation distortions in our reduced-form model by supposing that the additional price a_l creates a "distortionary impact" $k(a_l)$ that adds to the social cost of trades, with regularity conditions on $k(\cdot)$ to make our first-order approach below correct. To categorize possible exploitation distortions identified in the literature, we distinguish three extreme cases depending on exactly which trades are distorted: sophisticated-side distortions, naive-side distortions, and homogenous distortions. We consider each case in turn, solve for the market equilibria, and discuss applications that broadly fit into each of the cases. The type of the exploitation distortion will also be important for policy questions, such as the impact of third-degree price discrimination or consumer education discussed later.

(i) Sophisticated-side distortions. In this case, $k(a_l)$ arises only for trades with sophisticated consumers. Formally, we assume that sophisticated consumers' utility from purchasing product l is $v - f_l - k(a_l) - t|\gamma - l|$, while a naive consumer's utility is $v - f_l - a_l - t|\gamma - l|$, and firm l's cost of serving a consumer of either type is c. If a naive consumer anticipates taking the effort to avoid the additional price—but ends up not doing so—then her perceived utility is also $v - f_l - k(a_l) - t|\gamma - l|$. If a naive consumer does not think about the additional price at all, then her perceived utility is $v - f_l - t|\gamma - l|$. We analyze the former case, solving for the optimal contract that provides a perceived utility gross of transportation costs of $\hat{u}_l$ to consumers:

$$\begin{aligned} \max_{f_l, a_l} \quad & \alpha(f_l + a_l) + (1-\alpha)f_l - c \\ \text{s.t.} \quad & v - f_l - k(a_l) = \hat{u}_l. \end{aligned}$$

The constraint gives $f_l = v - k(a_l) - \hat{u}_l$. Plugging this f_l into the maximand and differentiating with respect to a_l yields that the equilibrium additional price $a(\alpha)$ satisfies

$$k'(a(\alpha)) = \alpha. \tag{1}$$

This trivial analysis already yields an economically important point: while a change in the level of competition (i.e., a change in t) does affect the equilibrium anticipated price (since it determines $\hat{u}_l$), it does not affect the equilibrium additional price. This contrasts with the prediction of many classical models, in which firms engaged in fierce competition charge prices close to marginal cost.

Since the equilibrium is symmetric, all consumers buy from the closest firm. Because the contract induces an exploitation cost of $k(a(\alpha))$ for every trade with sophisticated

consumers, the deadweight loss (DWL) relative to first-best—where consumers buy from the closest firm and $k(a) = 0$—is

$$DWL(\alpha) = (1 - \alpha)k(a(\alpha)). \tag{2}$$

Eq. (2) implies that an increase in the share of naive consumers, α, has an ambiguous effect on consumer welfare. It follows from Eq. (1) that an increase in α increases firms' incentive to focus their business model on the exploitation of naive consumers, increasing the additional price and thereby lowering the associated welfare of trading with sophisticated consumers. But with more consumers being naive, fewer consumers engage in inefficient avoidance behavior, increasing welfare.

A sophisticated-side distortion emerges in many papers due to a common source: that sophisticated but not naive consumers engage in costly behavior to avoid the additional price. In Gabaix and Laibson (2006) and Armstrong and Vickers (2012), firms charge high add-on prices (e.g., for room service in the case of a hotel or cartridges in the case of a printer) to profit from naive consumers who do not think about these prices. Facing such high prices, sophisticated consumers exert socially wasteful effort to avoid the add-on.[8] As a more subtle example, Seim et al. (2016) develop and empirically estimate a model of the Portuguese driving-school market in which firms provide a basic service—instruction up to the first driving exam—as well as additional services—instruction for consumers who need to repeat an exam. All consumers pay attention to the basic fee, but only sophisticated consumers pay attention to the additional fee. Higher additional fees therefore lead sophisticated but not naive consumers to exert more effort to avoid failing an exam. The authors provide empirical evidence suggesting that prices for the basic service fall in the number of competitors, but prices for the additional service do not, supporting the basic prediction of our simple model. Additionally, survey evidence indicates that a significant fraction of students are naive in the sense of being unaware of the additional fees, and that these students are more prone to overestimate their exam pass rate and less prone to engage in specific useful exam preparation techniques.

A variant of a sophisticated-side distortion arises in Grubb's (2015a) model of industries in which the price of a marginal unit depends on past purchases, and it is difficult for consumers to keep track of past usage. In his model, consumers need to pay an attention cost to recall past usage, and naive consumers underestimate their future attention costs. After having selected a contract, consumers can consume in up to two

[8] At the same time, Zenger (2013) points out that high add-on prices can be efficiency-enhancing if ex-ante avoidance is efficient, and there are (partially) naive consumers who underestimate but do not completely ignore their need for the add-on. These consumers may exert too little avoidance effort, and high add-on prices encourage them to exert more. Because in many applications the production cost of the add-on—and therefore the optimal ex-ante avoidance effort—is low, however, this consideration is often less important than the welfare loss for sophisticated consumers.

time periods, with their consumption values drawn independently in the two periods. In addition to a basic fixed fee, firms charge (potentially different) prices for consumption in the two periods, and possibly a penalty for consumption in both periods. A high penalty can be used to exploit naive consumers—who overestimate their probability of paying attention and avoiding the penalty—but it distorts sophisticated consumers' consumption decisions, and hence generates a sophisticated side-distortion.[9]

(ii) Naive-side distortions. In the case of a naive-side distortion, $k(a_l)$ arises only for trades with naive consumers. Formally, all consumers anticipate their utility from purchasing product l to be $v - f_l - t|y - l|$ and while sophisticated consumers forecast their utility correctly, a naive consumer's utility is actually $v - f_l - a_l - k(a_l) - t|y - l|$, and firm l's cost of serving a consumer of either type is c.[10] Solving for the optimal contract that provides a perceived utility gross of transportation costs of $\hat{u}_l$ to consumers:

$$\begin{aligned} \max_{f_l, a_l} \quad & \alpha(f_l + a_l) + (1 - \alpha) f_l - c \\ \text{s.t.} \quad & v - f_l = \hat{u}_l. \end{aligned}$$

Because the additional price increases profits and does not impact the constraint, the firm sets $a_l = a_{max}$. The exploitation of naive consumers is, thus, unaffected by their population share. Intuitively, because sophisticated consumers are not affected by the additional price, its choice does not affect a firm's market share, and hence firms choose that additional price that maximizes the profits earned from interacting with a naive consumer. In contrast to the case of a sophisticated-side distortion, the dead-weight loss in a symmetric equilibrium increases in the number of naive consumers; formally,

$$DWL(\alpha) = \alpha k(a_{max}).$$

While the literature has not focused on naive-side distortions, it is perhaps the most basic type of exploitation distortion that can emerge: the very fact that a naive consumer pays unexpected charges leads her (and not a sophisticated consumer) to miscalibrate her budget or intertemporal consumption. As another example, naive consumers who find out about unanticipated expenditures may get outraged and call the firm to complain, creating a cost for the firm and society. Furthermore, the higher the additional price,

[9] Grubb's (2015a) model, however, is not a perfect example of a sophisticated-side distortion. While absent sophisticated consumers the optimal contract induces efficient consumption by naive consumers, this need not hold when sophisticated consumers are present. In this sense, there may also be a naive-side distortion, albeit one would intuit that it is often less important.

[10] We suppose here that it is naive consumers who incur the exploitation cost. If instead the firm incurs the exploitation cost—as would be the case with legal and administrative costs of collecting additional prices—then a naive consumer's utility from purchasing product l is $v - f_l - a_l - t|y - l|$, firm l's cost of serving a sophisticated consumer is c, and its cost from serving a naive consumer is $c + k(a_l)$. A similar analysis shows that in this case the optimal additional price is implicitly defined by $k'(a(\alpha)) = 1$.

the more consumers complain, and hence the higher is the associated cost. A naive-side distortion also arises if firms offer products with useless but costly add-ons—e.g., rust proofing for a new car—that only naive consumers take.

(iii) Homogenous distortions. In the case of homogenous distortions, $k(a_l)$ arises in trades with both naive and sophisticated consumers. Formally, all consumers anticipate their utility from purchasing product l to be $v - f_l - t|\gamma - l|$, a naive consumer's utility is actually $v - f_l - a_l - t|\gamma - l|$, and firm l's cost of serving a consumer of either type is $c + k(a_l)$.[11] Proceeding along the same lines as in the case of a sophisticated-side distortion, the optimal contract in the homogenous-distortion case solves:

$$\begin{aligned} \max_{f_l, a_l} \quad & \alpha(f_l + a_l) + (1 - \alpha)f_l - c - k(a_l) \\ \text{s.t.} \quad & v - f_l = \hat{u}_l. \end{aligned}$$

Hence, the optimal additional price solves $k'(a(\alpha)) = \alpha$, and the dead-weight loss in a symmetric equilibrium equals

$$DWL(\alpha) = k(a(\alpha)). \tag{3}$$

In contrast to the case of sophisticated-side distortions, an increase in the share of naive consumers *unambiguously* lowers welfare.

There are several economically relevant examples of homogenous distortions. A homogenous distortion emerges in the credit-market model we analyze in detail in Section 2.5. In this model, naive time-inconsistent borrowers pay more interest than they expect, and to take advantage, firms overlend. Furthermore, because firms cannot distinguish naive and sophisticated borrowers, they overlend to everyone equally. A simple modification of the prime example of sophisticated-side distortions, where sophisticated but not naive consumers engage in costly avoidance of the additional price, also creates a homogenous distortion. In particular, if *both* naive and sophisticated consumers realize that they should avoid the additional price and take the same costly ex-ante steps to do so, but naive consumers pay more than they expect anyway, then the distortion affects all trades equally. Finally, firms may increase naive consumers' unexpected spending through product modifications, such as free alcohol and a glitzy environment to encourage gambling in a casino, that affect all consumers.

Intermediate cases. Many applications feature what might be described as an impure homogenous distortion: they involve both sophisticated-side and naive-side distortions, but these are not necessarily equal. At a basic level, a distorted price on an add-on typically distorts the consumption of both sophisticated and naive consumers, although

[11] As our analysis below highlights, the assumption that $k(a_l)$ is borne by the firm rather than the consumer does not affect the optimal additional price, firms' equilibrium profits, or consumers' equilibrium utility.

possibly to different extents. In Grubb (2009), for instance, cellphone consumers correctly predict their average demand for minutes, but they underestimate the variance in their demand. This mistake would not affect consumers if the price for minutes was linear. But to exploit the consumers' prediction error, a profit-maximizing firm charges a convex price, leading the consumer to underestimate how much she will pay in expectation. Because the marginal price per minute does not equal marginal cost, consumers' consumption decisions are distorted. Collecting information on mobile-plan choices and usage patterns, Grubb (2009) finds usage patterns consistent with the overconfidence explanation but not natural alternatives. Most notably, in a model of price discrimination in which high types also have highly variable demand, charging (more) convex prices for low types can discourage high types from taking the cheap package intended for low types. Yet in the data, the consumption of high types first-order stochastically dominates the consumption of low types.

To explain why the majority of customers who buy life insurance in the US do not hold the insurance until the end of the term, Gottlieb and Smetters (2012) propose a simple model in which consumers underappreciate non-mortality-type background risk, such as employment or health shocks, when buying life insurance. Because consumers underappreciate background risk, they underestimate the probability of lapsing that can occur after a bad background shock. Firms inefficiently front-load life-insurance premiums both to take advantage of unexpected lapsing—which leads consumers to forego cheap continued insurance—and to encourage further lapsing by depleting consumers' early resources. These inefficient loads can distort the lapsing decisions of both naive and sophisticated consumers. Gottlieb and Smetters carefully combine theoretical observations and empirical evidence to argue that alternative explanations do not provide a full account of the empirical patterns. Most notably, under a rational model of reclassification risk, front-loaded premiums discourage policyholders from lapsing after a favorable health shock, guaranteeing the integrity of the insurance pool. Yet health shocks are unlikely until older ages, so under reclassification risk younger consumers should not be paying significant loads—yet they pay the highest loads.

In Michel's (2016) model of extended warranties, naive consumers underestimate how costly it is to return a product, and hence overestimate the value of a warranty. Because naive consumers thereby misestimate the firm's warranty expenditure when offering a low-quality product, these consumers not only overestimate the value of the relatively useless warranty itself, but also the firm's incentive to produce high quality. As a result, firms have less of an incentive to produce high quality than they would with rational consumers, which can lead the firm to sell inefficiently low-quality products to naive consumers. And even for parameter values for which selling high-quality is still optimal, the firm distorts naive consumers' warranty terms in order to better exploit their return-cost misprediction. Due to screening issues (which we introduce in Section 3), however, the firm may also want to distort sophisticated consumers' contract offers, sometimes leading to distorted contracts for both types.

In the search model of Gamp and Krähmer (2017), firms choose between inefficient low quality and efficient high quality, and naive consumers erroneously believe that all firms offer high quality. The authors derive conditions under which sophisticated consumers always search for a high-quality firm, but naive consumers purchase immediately. Consequently, sophisticated consumers may inefficiently pay search costs, and naive consumers may obtain inefficiently low quality. As search frictions disappear, low-quality products come to dominate the market and naive consumers' purchases. Intuitively, the increase in competition resulting from the reduction in search frictions reduces the profit from offering a high-quality product, leading firms to focus their business model on exploiting naive consumers.

Externalities between naive and sophisticated consumers. To highlight a few additional issues in the case of homogenous distortions, suppose that the market is in the perfectly competitive limit ($t \to 0$). The utility of a sophisticated consumer is then

$$U_s(\alpha) \equiv v - f(\alpha) = v - c + \alpha a(\alpha) - k(a(\alpha)), \tag{4}$$

while the utility of a naive consumer is

$$U_n(\alpha) \equiv v - f(\alpha) - a(\alpha) = v - c - (1-\alpha)a(\alpha) - k(a(\alpha)). \tag{5}$$

An immediate implication of the above characterization is that, just as when there is no exploitation distortion, sophisticated consumers benefit from the presence of naive consumers. To see this, note that since $a(\alpha)$ maximizes $\alpha a - k(a)$, it must be the case that $\alpha a(\alpha) - k(a(\alpha))$ is strictly increasing in α.[12] Unlike in the case without exploitation distortions, however, whether naive consumers benefit from or are hurt by the presence of sophisticated consumers is in general ambiguous. Because competition forces firms to hand back the unanticipated payments from naive consumers, when more fellow consumers are naive, naive consumers get a larger fraction of this unanticipated payment handed back to them. But when more consumers are naive, firms also focus more on exploiting naive consumers, increasing the exploitation distortion borne by all consumers. Which effect dominates depends on the exact form of the exploitation cost function $k(\cdot)$. See Armstrong (2015) for a much more complete analysis of how naive and sophisticated consumers affect each other in equilibrium.

[12] This robust result does not depend on assuming a homogenous distortion but follows from the fact that naive consumers are more profitable. As a result, if in a competitive equilibrium the sophisticated consumers' utility would strictly decrease as α increases to α', a firm could deviate and offer the equilibrium contract for the case of α—which consumers strictly prefer—and earn positive profits as it would attract relatively more naive consumers.

Behavioral first welfare theorem. Another interesting property of the perfectly competitive outcome is that it satisfies a kind of "behavioral first welfare theorem:" it maximizes consumers' perceived utility (which equals $U_s(\alpha) \equiv v - f(\alpha)$) subject to a zero-profit constraint ($f(\alpha) = c + k(a(\alpha)) - \alpha a(\alpha)$)—that is, subject to the economy's resource constraints. The logic is in fact general to the current section's models: under perfect competition, consumers end up with a product they would be willing to choose if all products on which firms earn non-negative profits were available in the market. Firms are simply selling what consumers want to buy. If this was not the case, a firm could profitably offer a product consumers strictly prefer, undermining the logic of a fully competitive equilibrium.

This simple observation is related to a common question in seminars, and especially popular discussions: "whose fault is it" that consumers end up with suboptimal outcomes? We refrain from making moral judgments on the issue, but think that the above clarifies: firms do not necessarily have to be malicious in any sense for inefficient outcomes to occur. Indeed, even if firms have no explicit theory of consumer behavior, but are able to somehow find the profit-maximizing product, the same equilibrium obtains. Furthermore, even if firms do want to help consumers, in a competitive market they have little scope to do so. As pointed out in Appendix B of Gottlieb and Smetters (2012), the same market outcome remains an equilibrium outcome in a model in which one profit-maximizing firm competes with firms that attempt to maximize consumer welfare but must break even. These altruistic firms cannot offer consumers a deal that they are willing to accept and that yields higher welfare. We think of the outcome as due to the interaction between profit-maximizing firms with a deceptive technology and consumers' reaction to that technology.

We should note, however, that it will be more difficult to see firms in such an innocent light in Section 4, where they also manipulate consumer perceptions. In addition, when firms have market power, they do have scope for increasing consumer utility if they wanted to. As a case in point, Bubb and Kaufman (2013) develop a model of how mutuals—companies jointly owned by their customers—as well as nonprofit firms treat consumers who might be naive. They assume that managers of such companies have less financial incentives and incur a cost of exploiting customers, so that in equilibrium these firms charge (what we would call) a higher anticipated price and a lower additional price. Bubb and Kaufman (2013) offer evidence from the financial service market in the US that mutually owned firms offer lower penalties—e.g., default interest rates—and what they (reasonably) interpret as higher anticipated prices—such as higher introductory interest rates.

2.5 Example: deriving the additional price from primitives

So far, we have discussed the economic implications of hidden prices in a reduced-form model with an exogenously specified additional price. We now show for a simple ex-

ample how the additional price, and the distortion it generates, can be derived from a more basic model of consumer mistakes. Specifically, we sketch a simplified version of Heidhues and Kőszegi (2010)—which itself builds on the pioneering approaches of DellaVigna and Malmendier (2004) and Eliaz and Spiegler (2006)—to provide a micro-foundation for the additional price in a model of a credit market with partially naive time-inconsistent consumers.[13] We focus on two-part-tariffs, but the economic logic is the same with general contracts. Also, for simplicity we focus on a perfectly competitive market ($t = 0$).

In our three-period model, consumers may borrow money from lenders who have access to funds at zero interest and have no costs of making loan offers. In period 0, lenders make loan offers (b, r, d) that consist of a borrowed amount b, and interest rate r, and a discount d that—depending on the application—can be thought of as airline miles, cash back, or other credit-card perks. Upon observing the loan offers, consumers decide whether to accept a loan offer, and if so, which one. The utility of not accepting any offer is normalized to zero. Those consumers who borrowed b decide how much of their outstanding debt to repay in period 1. The remaining debt incurs interest and has to be repaid in period 2. So if $q \in [0, b]$ is the chosen repayment in period 1, the consumer needs to repay $(b - q)(1 + r)$ in period 2.

Crucially, consumers have a time-inconsistent taste for immediate gratification. Self 0's utility, which we take as relevant for welfare, is $u(b) - q - (1 + r)(b - q) + d$, where we assume that $u' > 0, u'' < 0$ and that $u'(0) > 1$ and $\lim_{b\to\infty} u'(b) < 1$. Self 0 hence trades off the benefit from borrowing (as well as the discount) with the *total* cost of repayment. Self 1, however, downweights period-2 repayment costs by a factor β satisfying $1/2 < \beta \leq 1$, choosing q to minimize $q + \beta(1 + r)(b - q)$. Self 0 has point beliefs $\hat{\beta}$ about her future β; that is, she believes that self 1 will choose q to minimize $q + \hat{\beta}(b - q)(1 + r)$. A consumer chooses a contract or the outside option to maximize her perceived utility, given her prediction about her own future behavior. Lenders know consumers' beliefs $\hat{\beta}$, and conditional on $\hat{\beta}$, there are two consumer types: sophisticated—who have $\beta = \hat{\beta}$—and naive—who have $\beta = \beta_n < \hat{\beta}$.[14] Note that since firms' and consumers' total utility is $u(b) - b$, social welfare depends only on the amount borrowed, and the efficient amount of borrowing satisfies $u'(b) = 1$.

We think of the firm as solving two interrelated problems: identifying the optimal contract that provides a given perceived utility $\hat{u}$ to consumers, and identifying the optimal $\hat{u}$. The main economic insights derive from the former part, so we solve this part. Observe that naive borrowers are willing to repay all of their loans in period 2 if $\beta_n(1 + r) \leq 1$, so that the highest interest rate at which they are willing to delay

[13] We developed this simplified version in Heidhues and Kőszegi (2017).

[14] The assumption that all consumers have the same beliefs at the time of contracting allows us to sidestep screening issues, including those we discuss in Section 3. Since all consumers have the same beliefs and preferences in period 0, they choose between contracts in the same way.

repayment is $r = (1-\beta_n)/\beta_n$. When setting this interest rate, the firm can collect interest of $(1-\beta_n)b/\beta_n$ from naive consumers. For this interest rate, $\hat{\beta}(1+r) > 1$, so all borrowers expect to repay their loans in period 1, and sophisticated consumers actually do. An optimal contract that generates unanticipated interest payments must solve

$$\max_{b,d} \underbrace{\left(b + \alpha(1-\beta_n)b/\beta_n\right)}_{\text{actual repayment}} - d - b$$
$$\text{subject to } u(b) - \underbrace{b}_{\text{expected repayment}} + d = \hat{u}.$$

Solving the constraint for d and plugging it into the maximand yields

$$\max_{b} \underbrace{\alpha(1-\beta_n)b/\beta_n}_{\text{unant. interest}} + \underbrace{u(b)-b}_{\text{social surplus}} - \hat{u}.$$

Hence,[15] the profit-maximizing loan size b satisfies $u'(b) = 1 - \alpha(1-\beta_n)/\beta_n$, and is therefore above the socially optimal level. Intuitively, the firm sets a high interest rate for delaying repayment that naive consumers do not expect to pay, but in the end do pay. These additional payments correspond to the additional price in our reduced-form model, and the amount consumers expect to pay to the anticipated price. Furthermore, to increase the amount of unexpected interest naive consumers pay, the firm induces overborrowing. Since all consumers get the same loan and hence all consumers overborrow, this distortion corresponds to a homogenous exploitation distortion in our reduced-form model. Nevertheless, sophisticated consumers benefit from the presence of naive consumers: in a perfectly competitive market, $d = \alpha r b$, so sophisticated consumers obtain credit below cost. In Heidhues and Kőszegi (2010), we argue that these properties have close parallels in real-life credit-card markets. Most credit cards do not charge interest on any purchases if a borrower pays the entire balance due within a one-month grace period, but deferring repayment to later carries large interest charges and potentially other fees.

A surprising feature of the equilibrium is that borrowing is discretely higher than optimal for *any* $\beta_n < \hat{\beta}$—that is, for an arbitrarily small amount of naivete. The consumer's small misprediction of her future preferences leads to a large welfare loss because the optimal contract hones in on and exacerbates her mistake: even though she mispredicts her future preference by only a little, she mispredicts her future behavior by a lot, and

[15] Observe that if a lender would set an interest rate at which sophisticated borrowers expect to delay repayment, this anticipated interest payment would feature in the participation constraint. As the rewritten objective function indicates, the lender earns lower profits when making loan offers that do not generate unanticipated interest.

with time inconsistency this has serious consequences. By honing in on and exacerbating the mistake and thereby collecting a high additional price, the firm can offer the most attractive-looking deal up front.

2.6 Identifying hidden prices from market data

In the course of our review, we discuss a variety of evidence indicating that in many markets firms are interacting with consumers who are not all rational. In a similar vein, we discuss how models of consumer mistakes can help explain a variety of evidence on market outcomes, and since most researchers also argue that it is difficult to find reasonable classical explanations for the same patterns, these patterns also constitute evidence of consumer mistakes. In all these instances, however, researchers have used special data opportunities, unique market features, or specific theoretical arguments to (directly or indirectly) document consumer mistakes. An important question we are often asked by regulatory agencies is whether there are any telltale signs that hidden prices are important in a market. Economic principles suggest a few possible "cook-book" approaches for regulators who have access to price and demand data. While researchers have used some of these approaches in individual cases, we are unaware of a synthesis in the literature, and provide some rough ideas here.[16] We hope future research will address this question more deeply and systematically.

Responsiveness to equivalent price changes. In all cases, we suppose that the product is potentially subject to two prices, f and a. If the price of a product is $f + a$, then with classical consumers the responsiveness of demand to an increase in f should be the same as the responsiveness of demand to an increase in a. Chetty et al. (2009) use this observation in the context of alcohol consumption in the US. Alcohol is subject to two kinds of taxes, an excise tax and a sales tax, and while the excise tax is included in the price posted in the store, the sales tax is added only at the register. Furthermore, these taxes vary considerably across states. Chetty et al. find a major violation of the above rationality principle: demand is much more responsive to excise taxes than to sales taxes.

Asymmetric demand responses. Suppose that f is the price for a base good and a is the per-unit price of an add-on to the base good. A consumer's type θ is drawn from the interval $[\underline{\theta}, \overline{\theta}]$ that admits a density $g(\theta)$, and the θ's are ordered such that for any prices, it is the consumers with higher types who buy the base good. A consumer's utility is quasilinear in money and the utility from the product and the add-ons, with her marginal utility of money being normalized to 1. Her outside option has a fixed utility

[16] We developed these preliminary ideas together with Takeshi Murooka, whom we want to especially thank for letting us use them in our survey.

level independent of prices in this market. Let $D_\theta(a)$ be the add-on demand of consumer θ *conditional on* purchasing the base good, and $V_\theta(f, a)$ the consumer's perceived indirect utility from purchasing the product when prices are (f, a). In addition, let $\tilde{D}_\theta(a)$ be the consumer's perceived add-on demand for a given a; for a rational consumer, we have $\tilde{D}_\theta(a) = D_\theta(a)$.[17] Finally, let $x_b(f, a)$ and $x_a(f, a)$ be the total base-good demand and add-on demand, respectively. Suppose we start from a given market situation (f, a), where the consumer who is indifferent between purchasing and not purchasing the base product has type θ_0.[18]

By the envelope theorem,

$$\frac{\partial V_\theta(f, a)}{\partial a} = -\tilde{D}_\theta(a) = \tilde{D}_\theta(a)\frac{\partial V_\theta(f, a)}{\partial f}. \tag{6}$$

Because small price changes only affect a marginal consumer's decision to purchase the base good, this immediately implies that

$$\frac{\frac{\partial x_b(f,a)}{\partial a}}{\frac{\partial x_b(f,a)}{\partial f}} = \tilde{D}_{\theta_0}(a). \tag{7}$$

Eq. (7) implies that the base-good demand of rational consumers ($\tilde{D}_\theta(a) = D_\theta(a)$) responds to the add-on price relative to the base-good price exactly in proportion to add-on demand. For instance, if the marginal printer consumer uses 40 cartridges for the printer, then a \$1 change in the cartridge price should have the same effect on printer demand as a \$40 change in the printer price—these have the same effect on the marginal consumer's total ownership costs. Shui and Ausubel's (2004) test of borrower rationality in the credit-card market tests exactly this prediction. They think of the introductory interest rate as the base-good price and of the post-introductory interest rate as the add-on price, and find that the above ratio is far below the rational level. Based on Eq. (7), the natural interpretation is that consumers underestimate their own add-on demand, i.e., long-term borrowing in the credit-card case.

Using that

$$\frac{\partial x_a(f, a)}{\partial f} = \frac{\partial x_b(f, a)}{\partial f} D_{\theta_0}(a),$$

[17] The perceived indirect utility of a consumer θ conditional on purchasing is $V_\theta(f, a) = U_\theta(\tilde{D}_\theta(a), 1) - f - a\tilde{D}_\theta(a)$, where $U_\theta(d, b)$ is the perceived direct utility of consuming d units of the add-on and $b \in \{0, 1\}$ units of the base good. We suppose that consumers' perceived preferences satisfy the usual regularity condition so that the perceived add-on demand is characterized through its marginal utility being equal to the add-on price.

[18] Thus, $x_b(f, a) = \int_{\theta_0}^{\bar{\theta}} g(\theta)d\theta$ and $x_a(f, a) = \int_{\theta_0}^{\bar{\theta}} D_\theta(a)g(\theta)d\theta$.

we can rewrite Eq. (7) in the following way:

$$\frac{\partial x_b(f,a)}{\partial a} = \frac{\partial x_a(f,a)}{\partial f} \cdot \frac{\tilde{D}_{\theta_0}(a)}{D_{\theta_0}(a)}. \tag{8}$$

For a rational consumer ($\tilde{D}_\theta(a) = D_\theta(a)$), Eq. (8) reduces to the well-known Slutsky equation. For consumers who might mispredict add-on demand, we think of Eq. (8) as the *misprediction-augmented Slutsky equation*: it says that the extent to which Slutsky symmetry is violated equals the extent to which marginal consumers underestimate add-on demand. As a result, testing Slutsky symmetry not only provides a test of consumer rationality, it also gives a quantitative estimate of consumers' degree of irrationality. We are unaware of empirical work using exactly this test.

The misprediction-augmented Slutsky equation has an interpretation beyond the setting of products with add-ons. Namely, the same logic applies to any setting in which a consumer purchases one good, and then has a chance to purchase a complementary good. Take, for example, cigarette consumption. Suppose that the initial decision of whether to buy the base good is whether to smoke a cigarette, and the demand function $D_\theta(a)$ represents the additional smoking if the consumer smokes the initial cigarette. Then, $\tilde{D}_{\theta_0}(a) < D_{\theta_0}(a)$ means that the consumer underestimates the addictiveness of smoking the current cigarette. In this reinterpretation of the model, Eq. (8) says that if the consumer is rational, the responsiveness of future cigarette demand to current price should be the same as the responsiveness of current cigarette consumption to the (net present value of) future price. Without emphasizing it, Becker et al. (1994) document a strong violation of this condition consistent with consumers' underestimation of the addictiveness of cigarettes.

Optimal price setting by firms. In addition to such demand-side tests, we can use a more structural approach, exploiting the assumption that firms set profit-maximizing prices. Suppose a profit-maximizing firm with marginal cost c of producing the add-on is setting (f, a). Consider an infinitesimal increase Δa in a combined with a decrease of $\tilde{D}_{\theta_0}(a)\Delta a$ of f. Since this leaves the marginal consumer indifferent, it does not affect the number of consumers who purchase the base good. For (f, a) to be profit-maximizing, we must therefore have[19]

$$\int_{\theta_0}^{\overline{\theta}} \left[(a-c)D'_\theta(a) + D_\theta(a)\right] g(\theta)\, d\theta - (1 - G(\theta_0))\tilde{D}_{\theta_0}(a) = 0.$$

[19] Below, we denote the cumulative distribution function of θ by G.

Rearranging gives

$$\frac{a-c}{a}\cdot\frac{a\int_{\theta_0}^{\overline{\theta}}D'_\theta(a)g(\theta)d\theta}{\int_{\theta_0}^{\overline{\theta}}D_\theta(a)g(\theta)d\theta}+1=\frac{\tilde{D}_{\theta_0}(a)}{D_{\theta_0}(a)}\cdot\frac{D_{\theta_0}(a)}{\int_{\theta_0}^{\overline{\theta}}D_\theta(a)g(\theta)d\theta/(1-G(\theta_0))}.$$

In more intuitive terms, we get the following expression for the optimal markup:

$$\begin{aligned}\frac{a-c}{a} &= \left(-\frac{1}{\text{elasticity of add-on demand}}\right)\\ &\cdot \left(1-\frac{\text{marginal consumer's perceived add-on demand}}{\text{marginal consumer's add-on demand}}\right)\\ &\cdot \left(\frac{\text{marginal consumer's add-on demand}}{\text{average add-on demand}}\right).\end{aligned}\tag{9}$$

Absent consumer mistakes, firms want to target price cuts to marginal consumers. Hence, if marginal consumers demand more of the add-on than do average consumers, then we would expect the add-on to be sold below cost. In contrast, if the marginal consumer has little add-on demand, then the add-on should be sold above cost. To see a potential application, consider again the credit-card industry. There, it is plausible that marginal consumers—being poorer—have a higher demand for credit than do inframarginal consumers, so if consumers were rational, credit-card interest rates should be below the cost of funds—a prediction that is drastically violated. As this example illustrates, while a precise demand-side test requires price responsiveness data, to predict the sign of the markup it often suffices to know the ratio between marginal and average add-on demand.

We want to emphasize that because consumer behavior and firm pricing can be driven by specific considerations outside our simple framework, the above screening tests should only be the start, and not the end of investigating a market. Following the failure of one of the above screening tests, a more in-depth theoretical and empirical analysis is needed to understand firm and consumer behavior.

3. PRICE DISCRIMINATION WITH NAIVE CONSUMERS

Price discrimination—treating different consumers differently—is an extensively studied topic in industrial organization. The classical approach typically presumes that the consumer heterogeneity motivating price discrimination pertains to preferences. In this section, we review a growing literature that studies price discrimination when consumer heterogeneity pertains to naivete. Such naivete-based discrimination raises theoretically novel issues. With rational consumers, a consumer's (potentially stochastic) behavior and her ex-ante beliefs about her behavior coincide. With naive consumers, they may

not, raising two new motives for discrimination. First, a firm may want to discriminate between consumers who have the same ex-post preferences (and hence behavior) but different ex-ante beliefs. Since consumers with different beliefs choose from available offers in different ways, it is possible to induce self-selection among them (i.e., screen them). This leads to second-degree price discrimination, which we discuss in Section 3.1.[20] Second, a firm may want to discriminate between consumers who have the same ex-ante beliefs but behave differently ex post. Since consumers with the same ex-ante beliefs (and ex-ante preferences) always choose from available options in the same way, it is impossible to induce self-selection among them. Hence, discrimination must be based on other information. This leads to third-degree price discrimination, which we study in Section 3.2.

A useful illustration of the above distinction is the contrast between environments in which sophisticated consumers can versus cannot avoid the additional price. If sophisticated consumers can avoid the additional price—as in the case of credit-card interest, for instance—then both naive and sophisticated consumers believe that they will not pay the additional price. The consumers therefore often choose from contracts in the same way, so it may not be possible to screen them. If sophisticated consumers cannot avoid the additional price—as with mutual-fund management fees all consumers pay—then consumers have different ex-ante beliefs, so it is possible to screen them. For instance, a naive but not sophisticated consumer is willing to take a small cut in the anticipated price along with a large increase in the additional price.

Although not conclusive, some evidence indicates that firms engage in both second-degree and third-degree naivete-based discrimination. Gurun et al. (2016) document that lenders targeted less sophisticated populations with ads for expensive mortgages. Ru and Schoar (2016) find that the offers credit-card companies send to less educated borrowers feature more back-loaded payments, including low introductory interest rates but high late fees, penalty interest rates, and over-the-limit fees. These patterns are consistent with third-degree price discrimination. In addition, Ru and Schoar also find that issuers attempt to screen consumers with menus of offers: cards combine rewards programs that appeal to less-sophisticated consumers with more back-loaded terms, and miles programs that appeal mainly to sophisticated consumers with more front-loaded fees.

3.1 Second-degree naivete-based discrimination

Naivete regarding time inconsistency. Starting with Eliaz and Spiegler (2006), a sizable theoretical literature studies screening potentially naive consumers according to their beliefs. Eliaz and Spiegler analyze perhaps the purest version of this issue, considering consumers who have the same ex-post preferences, but are differentially naive

[20] Our review of second-degree naivete-based discrimination has benefited from Kőszegi (2014).

regarding those preferences at the time of contracting. Specifically, a monopolist firm (or principal) offers a consumer (or agent) a contract that is contingent on an action the consumer will take later. The firm knows that the consumer's preferences will change by the time of the action, but a partially naive consumer assigns positive probability to her preferences not changing. Furthermore, the consumer is time-inconsistent, evaluating her action differently ex ante and ex post. This induces two effects. First, as in Section 5.2 below, the consumer prefers to commit her future behavior, and the firm wants to cater to this demand. Second, as in models of hidden pricing above, the firm wants to extract money from the consumer's misprediction of her own behavior—charging her a low fee if she does not change her mind and a high fee if she does. Eliaz and Spiegler show that—unlike in a classical screening model—near-sophisticated consumers are always pooled with sophisticated consumers and get the efficient contract that fully commits future behavior. Consider instead a contract that lowers the anticipated price and introduces a fee if the consumer changes her mind. This allows the firm to make a little money on near-sophisticated consumers' slightly wrong beliefs about paying the fee. The contract, however, is very attractive to much more naive consumers—who grossly underestimate the probability of paying the fee—increasing the information rent that must be paid to these consumers. Eliaz and Spiegler (2008) establish that the same result obtains if the consumer is not time inconsistent, but merely uncertain about her future preferences. Consumers who are sufficiently naive, however, get a contract with which they mispredict their future behavior.

In Heidhues and Kőszegi (2010), we analyze a model of a competitive credit market in which consumers differ in their beliefs about their time-inconsistency, $\hat{\beta}$, and for any $\hat{\beta}$ they may be partially naive or sophisticated. We solve the case in which firms can observe $\hat{\beta}$, and provide conditions under which consumers self-select into the same contracts even when firms cannot observe $\hat{\beta}$. Relatedly, in Heidhues and Kőszegi (2017, Appendix I.A), we analyze a variant of the credit-market model introduced in Section 2.5, and identify conditions under which all consumers borrow the same amount and pay the same interest as when $\hat{\beta}$ is observable to firms. Heidhues and Kőszegi (2010) illustrates the logic of self-selection through an example. Consider a consumer looking to buy a TV on sale financed using store credit that does not accrue interest for six months, but carries high interest and potentially penalties if she fails to pay back the loan within six months. A fancier TV is a better deal for two reasons: it comes with a larger discount and the larger loan's interest-free period is more valuable. To take advantage of the deal, however, the consumers needs to repay the full amount within six months; the larger the outstanding loan, the harder it is to do so. Hence, consumers choose TV's based on their beliefs about their ability to exert enough self-control to pay back the loan in time, sorting perfectly according to $\hat{\beta}$. This logic, however, requires assumptions on the distribution of types to guarantee that better TV's are more attractive for consumers repaying early and less attractive for consumers repaying late. We are

unaware of analyses under more general conditions, and more broadly do not know of further research on screening under competition in the types of models in Section 2 when sophisticated consumers can avoid the additional price.

General naivete. Johnen (2017b) compares monopolistic and competitive screening in a market where consumers can buy the product over multiple periods, and firms can automatically renew an uncanceled contract. Naive consumers overestimate their likelihood of canceling (for instance due to overconfidence regarding their limited memory or naivete regarding their future procrastination), and firms take advantage by automatically renewing and raising the price. Sophisticated consumers are attentive and correctly forecast their switching behavior, so the more the monopolist attempts to attract and then exploit naive consumers by a low price followed by a high price, the more sophisticated consumers can take advantage by signing up and then switching. Since—as in Eliaz and Spiegler (2006)—the aggressive exploitation of naive consumers implies that the monopolist must give sophisticated consumers a greater information rent, monopolists exploit naive consumers less than under full information. In contrast, under competition the full-information equilibrium—in which naive consumers are heavily exploited—remains an equilibrium whenever sophisticated consumers prefer their undistorted price offers to naive consumers' distorted offers. For these parameters, competition leads to a higher exploitation distortion than does monopoly.

Heidhues et al. (2017) model a competitive market in which firms can offer both a superior product—e.g., an index fund that is cheap to run—and an inferior product—e.g., a managed fund that generates more expenses than extra returns. Again thinking of management fees as the additional price, this is a situation where all consumers must pay the additional price. A natural separating equilibrium often arises: sophisticated consumers buy the superior product at a low (competitive) total price, but because naive consumers ignore the additional price, they buy the inferior product at a high total price. This is especially true if a floor on the front load—e.g., due to mutual-fund regulations discussed above—leads to positive profits on the inferior product.

Naivete about health and insurance. In the models we have discussed so far, a firm screens only with respect to consumers' beliefs. Yet in many economic settings, there are classical reasons for screening as well, and a natural question is how naivete-based discrimination affects outcomes in such settings. Perhaps the most important example is insurance. Sandroni and Squintani (2010) study a competitive insurance market with low-risk and high-risk consumers, some of whom are overconfident: they believe themselves to be low-risk when in fact they are high-risk. Because low-risk and overconfident consumers have the same beliefs, they cannot be screened, so whenever they buy, they must buy the same insurance contract. The price of this contract must in turn reflect the presence of some high-risk consumers, so that it is a bad deal (i.e.,

actuarially unfair) given consumers' beliefs. In contrast to the prediction of the classic insurance model of Rothschild and Stiglitz (1976), therefore, there may be a group of consumers who prefer not to buy any insurance.

Note that naive consumers have a qualitatively different effect on sophisticated consumers in Sandroni and Squintani's model than in our bare-bones model. In our main model, naive consumers cross-subsidize sophisticated ones, so that the presence of naive consumers benefits sophisticated consumers. Here, overconfident consumers render the low-risk insurance contract more expensive while leaving the price of the high-risk contract unchanged, hurting sophisticated consumers.[21]

Schumacher (2016) models an insurance market in which consumers—in contrast to the US but in line with other institutional settings—select a long-term health-insurance contract. Sophisticated consumers engage in a healthy lifestyle, and naive consumers believe that they will do the same, but they do not. With fixed and inflexible contracts, naive consumers exert a negative externality on sophisticated consumers because they select the same contract and generate higher expenses. But when firms can offer long-term flexible contracts that allow consumers to switch among insurance options, all consumers initially select partial insurance, and while sophisticated consumers stick to it, naive consumers switch to full insurance after realizing that they did not take care of themselves. Because naive consumers are locked in ex post, their firm extracts their extra benefit from switching to full insurance, increasing the ex post profits firms earn from naive consumers. As a result, the transfer from sophisticated to naive consumers is reduced, and sophisticated consumers may even benefit from the presence of naive consumers.

A basic testable implication of classical models of insurance markets with heterogeneity in risk, including Akerlof (1970) and Rothschild and Stiglitz (1976), is that higher-risk types purchase more extensive insurance. Contrary to this prediction, empirical research has often found no correlation or a negative correlation between risk and insurance coverage (see Chiappori and Salanie, 2000; Finkelstein and McGarry, 2006, for instance). Spinnewijn (2013) provides a way of accounting for these findings based on consumers' potentially naive and heterogenous beliefs in a model with both moral hazard and asymmetric information. In such a setting, consumers who are initially identical with respect to underlying health could have different beliefs regarding both the level of health risk they face and the sensitivity of their health risk to lifestyle. The insurance coverage a consumer purchases depends on the former belief, while her health behavior—and therefore also her resulting risk type—depends also on the latter. As a result, the correlation between insurance coverage and risk depends on the correlation between the two beliefs.

[21] See Armstrong (2015) for a detailed discussion of the effects naive and sophisticated consumers have on each other.

3.2 Third-degree naivete-based discrimination and privacy

We now turn to situations in which firms engage in third-degree naivete-based discrimination—they use outside information about naivete to discriminate between consumers. Because the primary way in which firms can obtain such information is by obtaining data on consumers from various sources, this issue is pertinent to the current concerns about privacy on the internet.

To motivate our discussion, we first argue (based loosely on Johnen, 2017a) that in competitive situations firms are more likely to engage in third-degree naivete-based discrimination than in classical third-degree price discrimination, so that understanding naivete-based discrimination is extremely important. In a competitive environment with classical consumers, information about an individual is of limited value, and hence firms are unlikely to seek it. Even with preference heterogeneity, if all goods are offered at marginal cost, a rational consumer simply self-selects and buys the product that matches her preferences best. The same is not true for naive and sophisticated consumers with the same initial beliefs. As we have mentioned, these consumers make identical choices at the contracting stage and hence cannot be screened. Furthermore, when sophisticated consumers can costlessly avoid paying the additional price and are therefore less profitable than naive consumers, or consumers differ in their profitability for another reason, firms have an incentive to identify naive consumers. And with firms' fast-growing ability to collect and analyze consumer data, it seems safe to assume that they respond to this incentive.

In Heidhues and Kőszegi (2017), we study naivete-based discrimination in a model similar to our bare-bones model above. To illustrate a key difference between naivete-based and classical preference-based price discrimination, the paper uses the example a monopolistic bank that sells bank accounts and overdraft services to consumers. Some of its consumers use the overdraft service heavily while others do so only rarely. Wanting to serve all customers as well as to benefit from heavy users' high demand for overdraft services, the bank charges a high overdraft fee and a relatively low account maintenance fee. The high overdraft fee, in turn, leads customers to overdraft less than what would be socially optimal. How will the bank react when it is able to identify the two customer groups? When engaging in preference-based price discrimination, a bank that gains the ability to discriminate low and heavy overdraft consumers, can now offer them different contracts. To increase its profits, the bank will lower the overdraft fee for heavy users and thereby increase social welfare, which it can extract through an increase in the maintenance fee.[22] In sharp contrast, when heavy users overdraft unexpectedly and by mistake, the bank will react to its new information by *increasing* the overdraft fee to heavy users, lowering social welfare. When heavy users overdraft unexpectedly, they are

[22] Similarly, not needing to worry anymore that attractive conditions for rare overdraft users will attract high overdraft users, the bank will lower the low users' overdraft fee also.

not willing to pay a much higher maintenance fee in exchange for a lower overdraft fee, as they do not anticipate benefiting from the lower overdraft fee. Hence, the only way to profit from these consumers is through a high overdraft fee.

More generally, we confirm that in both the monopolistic and imperfectly competitive cases, naivete-based discrimination is never Pareto-improving, and derive how the aggregate welfare impact depends on the type of exploitation distortion in the market. With homogenous distortions, naivete-based discrimination lowers welfare if the exploitation distortion $k(\cdot)$ satisfies a—we argue weak and empirically identifiable—specific condition. In the contracting setting of Section 2.5, for instance, a sufficient condition is that consumers' consumption-utility function $u(\cdot)$ satisfies prudence, which is a standard assumption and in line with empirical results. Intuitively, in that setting naivete-based discrimination leads firms to increase overlending to more naive borrowers and to decrease overlending to more sophisticated ones. But because increasing the distortionary overborrowing by a given amount is more harmful then decreasing it by the same amount is beneficial, the information will typically reduce total welfare.

In contrast, with sophisticated-side distortions perfect naivete-based discrimination always maximizes welfare: knowing that sophisticated consumers cannot be exploited, firms avoid setting an additional price when selling to a sophisticated consumer, eliminating any distortion; and when selling to a naive consumer, by assumption the additional price creates no distortion, so that total welfare is maximized.[23] Finally, in the case of naive-side distortions, naivete-based discrimination has no impact on welfare. Intuitively, because the additional price does not affect trade with sophisticated consumers, a firm maximizes the ex-post profits from naive consumers, leading it to offer the same additional price independently of what it knows about consumers' naivete.

Johnen (2017a) identifies a subtle source of firm profits due to third-degree price discrimination. To see his argument, take our bare-bones model of Section 2.3, and suppose that firm 0 (but not firm 1) can perfectly identify whether consumers in a given group are naive or sophisticated, and make different offers to these consumers. Such informational advantage could arise, for instance, if firm 0 has pre-existing customers whose behavior it can observe and analyze. Suppose also that products are homogenous ($t = 0$). Then, firm 1 cannot make a profitable offer below the average profitability of attracting a customer in that group, so that its anticipated price must be no less than $c - \alpha a_{max}$. Knowing this, firm 0 can set the anticipated price for naive consumers at (or slightly below) $c - \alpha a_{max}$, guaranteeing profits of $(1 - \alpha)a_{max} > 0$ per naive consumer.

[23] Kosfeld and Schüwer (2017) analyze a model in the Gabaix–Laibson tradition in which a firm receives a signal about a consumer's naivete after the consumer signs on, but before the bank sets the add-on price. Since this is a market with a sophisticated-side distortion, if the signal is perfect there is no inefficiency.

3.3 Other motives for discrimination

A few papers study price discrimination when naive consumers are present, but naivete is not the dimension along which firms discriminate. Grubb (2015a) asks how naivete affects classical preference-based discrimination. He considers services, such as mobile phones and bank overdraft protection, for which consumers may not know the marginal price of the next unit of service. If consumers correctly anticipate their probability of running into penalties, inattention to marginal prices actually facilitates efficient screening by firms. Intuitively, penalty fees for high usage prevent high-value consumers from taking the contracts offered to low-value consumers; yet because consumers do not know when they apply, these fees do not distort the consumption of low-value consumers.

Hoffmann et al. (2014) also study the effect of naivete on a type of preference-based discrimination. They consider a model in which each product has two dimensions, and firms—having obtained information about the consumer's preferences—reveal a consumer's utility in the dimension in which her utility is higher. Although such "selective disclosure" is biased, it still provides useful information, and hence raises consumer welfare unless consumers are naive about the bias, competition is limited, and firms are able to price discriminate.

To shed light on the role of privacy and data protection, Huck and Weizsäcker (2015) develop a model in which player A—say a consumer—interacts with player B—say a firm—and their transaction generates greater surplus if A shares more information with B. The problem, however, is that B may sell this information to player C, who can use the information in a way that harms A (for instance because C can price discriminate and extract more rent from A in another market). For any price that B may ask for the information, the types of C who buy are the ones who hurt A more (e.g., by being able to extract more rent from her). Huck and Weizsäcker allow for player A to be naive in the sense that she does not foresee the possibility of B selling information about her. They find that the presence of naive consumers makes the market for information more profitable, and hence naive consumers exert a negative externality on sophisticated consumers, whose data is also leaked more often. A regulatory intervention that makes it more difficult for B to sell information may make a partially naive A—who underestimates the harm selling the information to C can do to her—worse off by providing a false sense of safety.

4. PERCEPTION EXTERNALITIES

In this section, we discuss situations in which a firm's behavior affects not only how a consumer perceives the firm's offer, but also how she perceives alternative offers. In this sense, a firm exerts an externality on rivals through how and whether a consumer thinks about the rival's product.

Note that when generally interpreted, the classical industrial-organization topic of persuasive advertising—whereby advertising shifts a firm's demand curve—fits the above definition of competition with perception externalities. Persuasive advertising is probably based at least in part on psychological mechanisms.[24] Some of the models we discuss below can also be interpreted as being about persuasive advertising. But because much of the classical work is not based on psychological foundations, we do not review it here. For a review, see Bagwell (2007).

4.1 Educating, confusing, and manipulating

We first consider the possibility that a firm can take steps to help or hinder consumers in understanding the choices available in the market. We illustrate some of the main economic forces using a single framework that builds on that of Section 2, allowing us to sidestep mixed-strategy pricing equilibria. But we also comment on how the logic plays out in other models.

From an abstract point of view, most models in this subsection posit that a firm chooses not only price(s), but also an action that affects whether a consumer can find the best product available in the market.[25] Several papers investigate the theoretical aspects of such games. Piccione and Spiegler (2012) analyze a framework in which the probability that a consumer can compare products depends in a general way on firms' price-framing choices, and a consumer chooses a default option if she cannot compare products. Spiegler (2014) generalizes this framework by positing that firms' marketing messages induce a distribution of possible frames, and in combination with the firms' products, the frame determines the probabilities with which a consumer chooses each product. Spiegler (2016) provides a comprehensive review of this literature, especially focusing on the issues at play with obfuscation.

We incorporate specific examples of the above possibilities into the model of Section 2, continuing to focus on symmetric pure-strategy equilibria, and assuming to start that all consumers are naive. Specifically, we allow firms to manipulate whether consumers are sophisticated—i.e., whether they understand the additional price. Beyond

[24] Indeed, early writers in industrial organization had a rich—partly psychological—view of advertising that went beyond the information-provision view that dominated the second half of the 20th century. For example, Kaldor (1950) suggests that some advertising "attempts to influence the behavior of the consumer, not so much by enabling him to plan more intelligently through giving more information, but by forcing a small amount of information through its sheer prominence to the foreground of consciousness." When introducing the persuasive view of advertising Braithwaite (1928) already suggested that "the vast majority [...] [of advertising is] aimed almost exclusively at 'creating demand.' [...] Now advertising expenditure as thus defined aims at increasing sales by affecting the mind of the consumer. By various appeals it induces him to change his subjective valuation of the commodity." Along these lines Copeland writes that "[a]dvertising aims to educate consumers' tastes. Effective advertising is carried on in anticipation of a demand for itself which it produces" (Cherington and Copeland, 1925).

[25] The exceptions are the search-cost and dynamic models we discuss at the end of the subsection.

setting its anticipated price $f_l \geq \underline{f}$ and additional price $a_l \leq a_{max}$, firm l costlessly chooses whether to make the additional price obscure or transparent. We distinguish between two extreme forms of this revelation technology. Under an *education-favoring* technology, consumers become sophisticated if at least one firm chooses transparency. Such a technology is assumed in Gabaix and Laibson (2006) and much of the literature following it, where it is often labeled as unshrouding or educating the consumer. In contrast, under a *confusion-favoring* technology, consumers remain naive if at least one firm chooses price obscurity. A version of this assumption is made in Chioveanu and Zhou (2013), and other papers on strategic complexity. Both assumptions satisfy what Piccione and Spiegler (2012) call weighted regularity: each firm has a comparability choice that is imposed on the other firm whether or not the other firm likes it. Piccione and Spiegler show that under this assumption, the game-theoretic analysis of equilibrium is relatively simple.

As will be clear momentarily, outcomes can depend crucially on whether we are in an education-favoring or a confusion-favoring market. Unfortunately, however, we are unaware of any research on which assumption—or what mix of the assumptions—is more appropriate, and in what situation. This gap in the literature calls for both theoretical and empirical research. As argued also by Spiegler (2015), we need more structured theories of consumer naivete to help us understand what features or arguments lead consumers to understand products. And we need empirical work that sheds light on the determinants of consumer mistakes in different markets.

The question also arises what we mean by consumer education or confusion. From the perspective of the model's predictions, what matters is whether a consumer can choose the best product for herself, or chooses randomly; whether she fully understands the products or the incentives of firms is irrelevant. Hence, any aid that allows the consumer to choose the best product qualifies as education. Of course, this recognition raises another question: if a consumer does not fully understand the market, what types of communication regarding her best choice does she find credible? Once again, research on this issue is virtually non-existent.

Given that a firm's transparency choice is costless, a transparent market is always an equilibrium with an education-favoring technology, and an obscure market is always an equilibrium with a confusion-favoring technology. To see this, consider an education-favoring market, and suppose that all competitors of a firm are educating consumers. Then, all consumers will be sophisticated, whether or not the firm educates. Being indifferent, the firm is therefore willing to educate. This equilibrium, however, is arguably driven purely by the education technology. To shed light on more of the economics, we ask whether an obscure market is an equilibrium with an education-favoring technology, and whether a transparent market is an equilibrium with a confusion-favoring technology. This makes sense especially if transparency choices are costly, undermining the previous equilibrium based on indifference. Continuing with

the education-favoring technology, Heidhues et al. (2017) show that if an obscure market is an equilibrium when education is costless, then it becomes the unique equilibrium when education is costly, no matter how small the cost is.

Education-favoring technology. Consider the education-favoring technology, and take a candidate equilibrium with an obscure market in which $v > c + t$. As we have shown above, if the price floor is not binding ($\underline{f} \leq c + t - a_{max}$), then consumer mistakes have no effect on market outcomes. This immediately implies that no firm has an incentive to educate consumers: even if a firm did, it would not want to change its total price, and hence could not increase its profits. As a result, an obscure market is an equilibrium.

Eliaz and Spiegler (2011a, 2011b) arrive at a closely related conclusion using a different approach to consumer education. In their models, firms can influence what a consumer looks at—in marketing terminology, her consideration set—but conditional on what she sees, the consumer rationally applies a well-defined preference. This combination of assumptions captures marketing strategies that may draw consumers' attention, but do not provide information or directly affect preferences too much. A firm may, for instance, position a product in a way that it is more easily noticed, affect the consumer's online search, or remind the consumer of disadvantages of a product she already knows. More specifically, Eliaz and Spiegler (2011a) assumes that each of two firms chooses both a product and a marketing strategy, and prices are fixed. Consumers automatically consider the product of a randomly chosen firm, and depending on this default product as well as the rival's marketing strategy, they may also consider the rival's product. Eliaz and Spiegler (2011b) analyzes a variant of the same model in which each firm chooses a menu, and may include payoff-irrelevant items in a menu that attract consumer attention. For both models, Eliaz and Spiegler identify conditions under which there is an equilibrium that generates the same profits as those that obtain with rational consumers. Nevertheless, because persuading a consumer to consider a product is costly, there is a positive probability that a firm offers a socially suboptimal product and the rival does not persuade the consumer to abandon it. This means that the equilibrium is inefficient, and because firms' profits are not affected, all of the inefficiency due to consumers' bounded rationality is borne by consumers. These results are reminiscent to those above, as well as several situations we have discussed in Section 2, where consumer naivete does not necessarily have implications for firms' profits, but could nevertheless lower consumer welfare. Unlike above, however, in Eliaz and Spiegler marketing is used with positive probability. Furthermore, Eliaz and Spiegler establish that any equilibrium with rational-consumer profits features what they call an "effective marketing property:" if a firm draws a consumer's attention purely because of a costly marketing message or attention grabber, then it ends up selling to the consumer.

The analysis is different if charging an additional price is subject to an exploitation distortion. Suppose that charging a leads to a homogenous distortion $k(a)$ paid by the

firm. In an obscure market, it must be the case that each firm chooses $a(1)$ satisfying $k'(a(1)) = 1$, and the firm makes a profit of $a(1) - k(a(1))$ from the additional price. In a candidate equilibrium, the anticipated price then becomes $f = c + t - (a(1) - k(a(1))$. But this is not an equilibrium: a firm can deviate by educating consumers, charging consumers $f = c + t + k(a(1)), a = 0$, attracting the same number of consumers while earning higher margins. Intuitively, the firm offers an efficient pricing scheme instead of an inefficient one, informs consumers of this, and captures as profits the eliminated deadweight loss. This formalizes a verbal argument by Shapiro (1995), who implicitly assumes an education-favoring technology, and concludes that a market in which the obscure nature of prices generates an inefficiency could never arise in equilibrium.

Now suppose that the price floor is binding ($\underline{f} > c + t - a_{max}$). Then, in the candidate obscure equilibrium, firms' margins are high, and therefore each firm wants to attract consumers to increase profits. But since the anticipated price is already at the floor, this is only possible by educating consumers and lowering the additional price. Would a firm want to do so?

If the total price is lower than consumers' value ($\underline{f} + a_{max} \leq v$), then the answer is clearly yes. Intuitively, since consumers value the product highly, a firm can attract more consumers by educating them about high prices, but at the same time lowering prices a little bit. This insight is a variant of a general result by Piccione and Spiegler (2012, Proposition 1) derived in a homogenous good framework: that if the outside option is irrelevant for consumers' choices (in Piccione and Spiegler's case because firms are restricted to price below consumers' value) and a firm can educate consumers, then the Bertrand outcome obtains. In this case, a firm offering a better deal always wants to make sure consumers understand this, so that firms cannot escape the Bertrand logic.

If the total price is greater than consumers' value ($\underline{f} + a_{max} > v$), however, the logic of the interaction is different. Then, as explained by Heidhues et al. (2017), if a firm educates consumers and cuts its additional price by a little bit, consumers' realization that the price is so high leads them not to buy. Thus, the firm can attract consumers only if it cuts the additional price by a discrete margin. Since this may be unprofitable, the firm may prefer not to educate.[26]

Gabaix and Laibson (2006) identify another reason that firms may not want to educate consumers. As we have explained in Section 2.3, in their model sophisticated consumers inefficiently avoid the additional price, but benefit from naive consumers through a cross-subsidy. A firm can educate consumers and propose a low additional price that sophisticated consumers prefer not to avoid, increasing efficiency. A sophisticated consumer, however, might prefer to trade inefficiently with another firm and

[26] Formally, when all firms charge prices $\underline{f}, a_{max}$, each firm earns $(\underline{f} + a_{max} - c)/2$. A firm that educates consumers can charge a total price of at most v. Even if it captures all consumers, it earns a profit of $v - c$, which may well be less than $(\underline{f} + a_{max} - c)/2$. If this is the case, the firm does not want to educate.

obtain the cross-subsidized price, rather than trade efficiently and receive no cross-subsidy. As a result, education may not be a profitable marketing strategy. In particular, because the cross-subsidy the sophisticated consumers like is increasing, and the gain from trading efficiently is decreasing, in the number of naive consumers, education is less likely when the proportion of naive consumers is high. And because the efficiency gain is increasing in sophisticated consumers' avoidance cost, an increase in the avoidance cost makes education more likely.

Heidhues et al. (2017) explore how the decision of whether to educate consumers interacts with features of the market and the product being sold. As they highlight, deception is especially stable for bad products that would not survive in the market absent deception. This is easiest to see when $v < c$ in our basic framework above; in that case, a firm that educates consumers can charge at most v, which is unprofitable. As a result, a deceptive equilibrium always exists. The result extends to a multi-product setting with a socially inferior and superior product. If the superior product is sold competitively, then whenever a firm educates, consumers will purchase the superior product and hence a firm selling the inferior product will never benefit from educating. And because the superior product is sold competitively, firms do not earn a positive margin when doing so, implying that even firms selling the superior product have no incentive to educate. In the presence of a price floor, then, firms earn positive profits not despite selling an inferior product but because they sell an inferior product. And the addition of a superior product can expand the scope for profitable deception by reducing the incentives to educate. As a potential example, firms have limited incentive to educate consumers about the inferior nature of managed mutual funds, because consumers would then invest in index funds with very low margins.

More generally, the effect of competition on education depends in complex ways on market specifics. Heidhues et al. establish that if the floor on the anticipated price is binding and there is a single socially valuable good ($v > c$), then with sufficiently many firms in the market an obscure equilibrium ceases to exist. With many firms vying for the obscure market's profits, there must be a firm whose market share and hence profits are very small. This firm can educate consumers and cut the total price to v, attracting all consumers and earning positive profits bounded away from zero. In such environments, competition-policy measures that increase the number of firms help facilitate market transparency. In contrast, because competition in the market for a superior product lowers the incentive to educate consumers about the inferiority of an alternative product, competition is not uniformly beneficial.

Wenzel (2014) analyzes the effect of competition in a variant of the Gabaix and Laibson (2006) setup in which the share of naive consumers who become educated is increasing in the number of firms who educate, but does not jump to one once a single firm educates. This pattern is plausible in many settings: consumers may overlook an education attempt or not take it seriously until they see the same warning again and

again. Wenzel argues that—in contrast to the model of Gabaix and Laibson (2006)—consumer education is more likely to occur in a more competitive market. For starters, a transparent equilibrium is more likely to exist with more firms. Intuitively, the more rivals educate, the less likely it is that a consumer is still naive, and hence the less profitable it is to try to exploit naive consumers at the cost of cross-subsidizing sophisticated consumers. At the same time, an obscure equilibrium is also more likely to exist with more firms. Nevertheless, Wenzel (2014) argues that with many firms reasonable equilibrium selection in the spirit of risk dominance tends to favor a transparent equilibrium. Intuitively, shrouding is risky in that the pricing can be exploited by educated consumers, and with more competitors the risk that someone may decide to educate consumers is weighted higher. A firm that educates does not face such risk.

Murooka (2015) investigates whether commission-motivated intermediaries—such as mortgage brokers, financial advisors, or insurance salespeople—can be relied upon to educate consumers about hidden fees. Murooka assumes that a transparent firm (which charges no additional price) and a deceptive firm (which charges an additional price) compete by choosing prices for consumers and commissions for intermediaries. Upon learning the firms' offers, competing intermediaries decide which product to offer to consumers and whether to educate consumers about the deceptive firm's additional price. Murooka shows that an obscure equilibrium exists if and only if the additional price is large. This means that intermediaries fail to fulfill their role of educating consumers exactly when that role is most important. Intuitively, when the additional price is large, the deceptive firm can afford to pay a large commission to intermediaries, who then prefer to sell the deceptive product to a few consumers rather than attract many consumers by educating. Because intermediaries need to be bribed not to sell the transparent product, they earn supra-competitive commissions. And because the commissions are ultimately paid by consumers, consumers are worse off when intermediaries can educate than when they cannot.

Confusion-favoring technology. Now suppose that we are in a confusion-favoring market. Then, a transparent market is an equilibrium if there is no exploitation distortion and the market is sufficiently uncompetitive—i.e., t is sufficiently high—for the price floor not to be binding ($\underline{f} \leq c + t - a_{max}$). In particular, if firms set $f = c + t - a_{max}$, $a = a_{max}$, then there is nothing to be gained from confusing consumers: confused consumers choose the firm with the lowest anticipated price, and a firm cannot profitably lower the anticipated price below $c + t - a_{max}$.

If either of the conditions fails, however, a transparent market ceases to exist. Suppose that there is a homogenous distortion paid by the firm with $k(a) > 0$ for any $a > 0$, and the price floor is not binding. Then in a transparent market firms choose $f = c + t$, $a = 0$, avoiding the distortion. But this is not an equilibrium: a firm can profitably confuse consumers and increase a. Unfortunately, therefore, in equilibrium firms confuse consumers exactly when that leads to a distortion.

Alternatively, suppose that there is no exploitation distortion, but the market is sufficiently competitive—i.e., t is sufficiently low—for the price floor to be binding ($\underline{f} > c + t - a_{max}$). In a candidate transparent equilibrium, firms set a total price of $c + t$, and each gets half of the market. A firm can therefore gain by confusing consumers and setting $\underline{f}, a_{max}$: because $\underline{f}$ is the lowest possible anticipated price, the firm still gets at least half of the market, and because $\underline{f} + a_{max} > c + t$, it has a higher margin. Hence, the obscure market is the unique equilibrium. This insight captures a general comparative static obtained by Carlin (2009) and Chioveanu and Zhou (2013) (which we discuss in more detail below): that as a market becomes competitive, it also becomes obscure. Intuitively, in a less competitive market firms can maintain high margins by virtue of their market power. In a competitive market, however, high margins require obfuscation, tending to shift the nature of the market towards an obscure one. Carlin argues that this prediction is roughly consistent with evidence by Hortaçsu and Syverson (2004) that in the 1990s S&P 500 index funds experienced both significant entry and an increase in fees.

The effect of competition can therefore be radically different with education-favoring and confusion-favoring technologies: in a very competitive market with a binding price floor and high consumer value ($v \geq \underline{f} + a_{max}$), a transparent market is the unique equilibrium with an education-favoring technology, and an obscure market is the unique equilibrium with a confusion-favoring technology. But the situation is not symmetric. To see this, consider a firm that can unilaterally affect (either educate or confuse) a small number of consumers, and cannot price discriminate. In an obscure market, the firm would not benefit much from educating a few consumers: to attract a significant fraction of them, it would have to lower its margin also on the many consumers who remain confused. In a transparent market, however, the firm would want to confuse even a few consumers because it could then sell to these consumers at a high margin. These insights imply that so long as price discrimination is infeasible, economic forces favor obfuscation rather than education in a competitive market, and obfuscation can dominate outcomes even if the scope for unilateral obfuscation is quite small. Indeed, in the models of Carlin (2009) and Chioveanu and Zhou (2013) partial education and partial obfuscation are equally feasible, and both papers find that competition leads to obfuscation. Our logic indicates that if price discrimination is feasible, then there is more scope for partial education—a firm could educate a few consumers, and sell at a cheaper price only to them—but this question has not been investigated so far.

If it is costly to confuse consumers, then obfuscation is of course less likely to happen, especially when there is competition. In Carlin and Manso's (2011) reduced-form model, obfuscation carries a fixed cost, so competition decreases obfuscation because it lowers each firm's benefit from increasing the share of confused consumers. Glaeser and Ujhelyi (2010) develop a model in which firms compete in a Cournot fashion, and each firm can increase industry demand through misleading claims that induce consumers to

overvalue the product. Glaeser and Ujhelyi argue that their model captures, for instance, false health claims regarding ineffective and dangerous medications, as well as (hidden) suggestions to the health benefits of smoking, both of which used to be common. In this framework, some of the benefit of increasing industry demand accrues to other firms, so there is a free-rider problem among firms. Because the free-rider problem is greater with many firms, misinformation decreases in the level of competition.[27]

Other approaches to obfuscation. Obfuscation is defined above as hiding the additional price, but there are other natural ways to define the concept. A number of papers think of obfuscation as making it difficult to compare prices. A firm could, for instance, break prices into multiple components in an economically meaningless way, or condition the price on unnecessarily many contingencies, to make the total price unclear. In Carlin (2009), each firm makes a pricing choice as well as a complexity choice, and firms' complexity choices jointly determine the share of consumers who are informed—and therefore choose the cheapest firm—or uninformed—and therefore choose randomly. This creates a Varian-type (1980) search model in which a mixed-strategy pricing equilibrium obtains. As might be expected, a firm that chooses a relatively low price aims for low complexity to be able to attract consumers, and a firm that chooses a relatively high price aims for high complexity to hide how expensive it is. Carlin finds that as the number of firms increases, the probability that a firm chooses high complexity increases. In Chioveanu and Zhou (2013), firms can adopt simple or complex frames, and consumers are less likely to be able to compare two firms' prices if the firms adapt different frames or they both use complex frames. A consumer chooses a firm with positive probability if no competitor to which she can compare the firm's price has a lower price. Again, an increase in the number of firms leads firms to use complex frames more often.

Other papers approach the question of obfuscation from the perspective of search costs. Just like a firm might want to make it more difficult for a consumer to compare prices in the market, a firm might want to increase consumers' costs of searching prices in the market. More interestingly, Ellison and Wolitzky (2012) and Wilson (2010) identify reasons why a firm might unilaterally want to increase consumers' costs of searching only *its own* price—for instance, by removing itself from a price-comparison website. Ellison and Wolitzky show that this can benefit the firm if search costs are convex, so that increasing the cost of searching the firm's price also increases the cost of searching further. And Wilson shows that the same can benefit a firm by softening price compe-

[27] Observe, however, that from a positive perspective, misinformation in the model works much like a product innovation. In richer models of competition with product specific misinformations, classic results in the innovation literature suggest that competition may increase or decrease firms' misinformation incentives.

tition and thereby increasing the profits from consumers who find the firm despite the higher search cost.

At their core, all of the above papers on obfuscation or the lack of education depend on an insight already recognized by Scitovsky (1950): that consumer ignorance is a source of oligopoly power, and hence firms are often interested in creating ignorance.[28] There is some empirical evidence consistent with this basic insight. Ellison and Ellison (2009) study an online computer-parts retailer that gets most of its business from a price-comparison website. Firms quote prices for a low-quality product on the price-comparison website, and indeed consumers' elasticity of demand with respect to this price is extremely high. But a retailer can charge for upgrades, shipping, and other add-ons once the consumer is referred to its site, leading to higher markups than would be implied by the elasticity of demand. These results suggest that a kind of floor on the low-quality product's price is in operation. In fact, Ellison and Ellison (2009) provide direct evidence of the adverse-selection-based foundation for the price floor (Section 2.2) by documenting that a lower rank on the comparison site is associated with a lower proportion of consumers buying upgrades. The observation that the rank affects what the consumer buys also suggests that consumers' understanding of the market is meaningfully limited.[29]

Hastings et al. (2017) study a completely different setting, Mexico's privatized market for social security. Being heavily regulated, funds' investment strategies are essentially homogenous. Furthermore, concentration in the industry is relatively low. While one would therefore expect funds to compete aggressively on price, prices in this market are shockingly high. For instance, a 100-peso deposit by a typical Mexican worker into an account that earned a five percent annual real return would be worth only 95.4 pesos after 5 years. Hastings et al. document that a major reason for the high prices was the low price sensitivity of investors, especially when investors were exposed to a large sales force. These findings are consistent with the models of obfuscation we have discussed above, as well as with the central notion in Murooka (2015) that sales agents direct consumers toward expensive products.

[28] Bar-Gill (2009) makes an observation in the same spirit. He argues that the complexity of the fees lenders could impose in the subprime mortgage market rendered it exceedingly difficult to compare products, so—despite the seemingly competitive nature of the market by conventional measures of concentration—lenders acted as local monopolies.

[29] The last conclusion is confounded by the possibility that the firm's rank on the price-comparison website is correlated with the value of its other offerings. Ellison and Ellison's dataset allows them to conduct an alternative test avoiding this confound. The retailer from which they have data operates two websites with identical products. If consumers fully understood all prices, then conditioning on a consumer's decision to purchase from one of the two retail sites, the retail sites' ranking on the comparison site should not predict which retail site the consumer purchases from. This prediction is violated in a major way.

The models we have reviewed in this section are all static in the sense that firms' decisions to manipulate consumer understanding occur simultaneously with their pricing decisions. Realistically, however, manipulations—e.g., education—are often longer-run campaigns. There is little research on the implications of this possibility. As one exception, Dahremöller (2013) extends and modifies Gabaix and Laibson's (2006) model by considering a duopoly setting in which firms make their observable education decisions before their pricing decisions, and their costs of producing the add-on are different. Then, outside a knife-edge case either the efficient or the inefficient firm wants to educate to disproportionately hurt the other firm's add-on profits, making the competitor less aggressive in the base-good market. In particular, the efficient firm may want to educate if this leads to a lower add-on price and hence higher add-on demand; and the inefficient firm may want to educate if this lowers the efficient firm's add-on sales.

Building on the general formulation of product comparability in Spiegler (2016), de Roos (2017) investigates how limited comparability affects the ability to collude in a homogenous-good industry. Obfuscation implies that a deviator attracts fewer consumers, and therefore benefits less, from marginally undercutting the cartel price; but obfuscation also makes it more difficult to punish a deviator by attracting consumers away. Despite the opposing forces, for many (but not all) types of obfuscation there is a tendency for obfuscation to aid collusion. Intuitively, while the deviator cannot systematically attract consumers away from all competitors, rivals can punish the deviator so long as consumers consider at least one more product.

4.2 Endogenously determined attention

In this section, we discuss a body of research studying how firms behave when they do not have actions to explicitly manipulate consumer perceptions at their disposal, but nevertheless consumer attention depends endogenously on the market environment. The papers use a variety of modeling assumptions regarding whether and how consumers compare products, but there are a few common threads.

Salience. The largest set of papers builds on the model of salient thinking by Bordalo et al. (2013).[30] The central aspect of the model is that the way in which a consumer trades off between attributes of a product (e.g., quality and price) depends in specific ways on the choice set facing the consumer. In a property called ordering, an attribute of a good is more salient, and hence gets a larger weight in choice, if it is further from the average value of the attribute in the choice set. And in a property called diminishing sensitivity, the salience of a product's attribute decreases if all products' values in that

[30] See Herweg et al. (2017) for a more detailed review of the model's implications for consumer behavior, and the resulting effects on markets, than we have scope for here. Our summary of the literature has benefited tremendously from their discussion.

attribute increase by the same amount. We use these properties to explain the market implications of salience.

In Bordalo et al. (2016), two firms first set their quality levels, and then choose prices. A firm's unit cost of production has a quality-independent component and a quality-dependent component. If the fixed component is relatively high, then prices will be high relative to quality, so diminishing sensitivity implies that price differences will not be as salient to consumers. As a result, a quality-salient, "de-commoditized" equilibrium obtains, in which quality provision is inefficiently high. In contrast, if the fixed component of the cost is relatively low, then prices will be low relative to quality, so price differences will be more salient to consumers. As a result, a price-salient, "commoditized" equilibrium obtains, in which quality provision is inefficiently low. Bordalo et al. argue that these results can help understand shifts in the nature of markets, such as the drastic transformation of the US coffee market in the 1990s.

Helfrich and Herweg (2017) and Dertwinkel-Kalt and Köster (2018) both identify salience-based reasons for manufacturers of quality products to ban online sales of their products. The mechanisms driving this preference, however, are completely different in the two papers. In Helfrich and Herweg, competition from online sales lowers prices in the market, increasing the salience of prices due to diminishing sensitivity and hence decreasing the markup that the quality manufacturer can charge. In Dertwinkel-Kalt and Köster, lower online prices render the product's price more salient due to ordering, hurting the manufacturer.

Inderst and Obradovits (2016) ask how salience affects the logic of markets with hidden prices. The greater is the price component that firms can hide, the lower are anticipated prices. Due to diminishing sensitivity, this increases the price sensitivity of consumers, and makes a price-salient equilibrium more likely, resulting in inefficiently low-quality products being sold and produced. As greater competition lowers headline prices, it exacerbates this adverse effect. Inderst and Obradovits (2017) exploits a similar logic in the context of loss leaders. If a product is used by retailers as a loss leader, then its price will be low, rendering price decreases more salient to consumers. By implication, retailers may prefer to sell loss leaders of inefficiently low quality, justifying the view that loss leading results in a race to the bottom in product quality.

Apffelstaedt and Mechtenberg (2017) analyze retailers' equilibrium product lines when consumers are partially naive regarding their sensitivity to in-store context effects. The model allows for a range of possible context effects, including those in Bordalo et al. (2013), Kőszegi and Szeidl (2013), and Bushong et al. (2017). A consumer first chooses among retailers, not (fully) anticipating that she will be subject to context effects later on. Once she is at her chosen retailer, she is locked in, and her decision depends on the products available at that retailer. Similarly to Heidhues and Kőszegi (2010), a consumer is attracted to a retailer by a bait product that she thinks she will buy, but once at the store she is induced to switch to a different, more profitable product. Unlike in Heidhues and

Kőszegi (2010), the more profitable product may not only be more expensive than the bait ("up-selling"), but also less expensive than the bait ("down-selling"). Furthermore, the retailer also uses a third, decoy product to manipulate the consumer's preferences to induce switching. Salant and Siegel (2017) develop a related model in which consumers are subject to framing effects at the store but not outside the store, and return an item if it is below their value ex post. Up-selling is also possible in their model, and allows a firm to overcome consumer-protection regulations. A firm may also use framing to lower screening distortions when there are heterogenous consumers, possibly increasing profits as well as consumer welfare.

Limited attention. Several papers explore implications of the idea that the attention consumers can devote to market information is limited, and a consumer may make strategic decisions in how to allocate her limited attention. That consumers use heuristics instead of fully attending to all available information, and market prices respond to this fact, is documented by Lacetera et al. (2010) and Busse et al. (2013). Lacetera et al. find that the price of used cars exhibits discontinuous drops at multiples of 10,000-mile odometer readings, and that this is driven by limited attention by final consumers rather than wholesalers. Studying German data, Englmaier et al. (2018) observe similar price drops at multiples of 10,000-kilometer odometer readings. In addition, they find a large increase in a used car's price if it is registered in January rather than December of the previous year—presumably because consumers pay limited attention to the initial registration date as well. That market participants also allocate their attention strategically is documented by Bartoš et al. (2016). They find that negatively stereotyped minority names on applications reduce employers', but increase landlords' effort to obtain more information about the applicant. This is consistent with a model of statistical discrimination in which employers' goal is to seek out the best candidates—leading them to immediately reject negatively stereotyped applicants—and landlords' goal is to screen out the worst tenants—giving them reason to inspect negatively stereotyped applicants more closely.

Martin (2017) analyzes the informativeness of prices in a standard pricing game when buyers are rationally inattentive to quality in the sense of Sims (2003) and the literature following it. A seller, knowing whether its product is of high or low quality, decides between a high and a low price. After observing the price, the buyer—not wanting to buy the low-quality product at the high price—chooses the extent to which she attends to the quality of the product. In the hope of making a large profit, the low-quality seller sets a high price with positive probability, lowering the informativeness of prices. Often, a decrease in the cost of attention lowers the probability that the low-quality seller sets a high price, but in some situations the opposite is the case. In particular, if the outside option is ex-ante superior, then a decrease in the cost of attention can increase the probability that the buyer abandons her outside option and purchases, potentially increasing the low-quality seller's incentive to mimic the high-quality seller.

When consumers can only compare a limited number of products, Hefti (2016) shows that even as more and more information becomes available and more and more products can be found (e.g., online), prices do not converge to marginal cost as entry costs approach zero. To understand the key intuition, consider the Salop (1979) model of competition on a circle, and suppose that a consumer has the capacity to look at two different randomly chosen products in the market. If there are two firms, the consumer can check both, leaving her on average 1/4 away from her ideal variety and leading to standard Salop-style competition. If instead there are four firms, then the consumer is likely to choose neighboring firms, leading her to be further than 1/4 from her ideal variety and also softening price competition. Diversity therefore does not benefit consumers who cannot sort through it. Hefti and Liu (2016) arrive at a similarly dim view of targeted advertising—advertising aimed at consumers likely to benefit most from the firm's product. Consider again a symmetric Hotelling model with two firms, and suppose that advertising costs are low. If consumers have unlimited attention, a firm just targets consumers for whom it is the best match, as it cannot get the other firm's consumers. But if consumers have limited attention—they may not observe the firm's advertising—then a firm targets everyone, hoping to defend its own consumers from the rival and also hoping to attract the rival's consumers. Hence, even though the technology exists for targeted advertising, firms engage only in mass advertising.

De Clippel et al. (2014) study a different form of competition with strategically inattentive consumers. Consumers observe the price of the market leader in each of multiple markets, and can also inspect competitors' prices in a given number of markets of their choice. By lowering its price, a market leader increases the chance that the consumer ignores the rival and buys from it, so that leaders effectively compete for consumer *inattention* across markets. An increase in consumers' capacity to inspect markets can induce leaders to focus on exploiting the most inattentive consumers, lowering competition and increasing prices.

In Heidhues et al. (2018), products have multiple price or quality components (such as a base price and an add-on price or a price and a level of safety), and consumers can only evaluate a limited number of relevant components. This means that if a consumer carefully studies many products, she has less attention left for comparison shopping. Heidhues et al. show that because low-value consumers are often more likely to study—and therefore less likely to comparison shop—than high-value consumers, the average price consumers pay can be increasing in the share of low-value consumers. This prediction helps explain why a number of essential products are more expensive in lower-income neighborhoods.

5. RESPONDING TO CONSUMER PREFERENCES

In much of classical industrial organization, a consumer's preferences are extremely simple—she has a fixed valuation for each product, and may also be subject to search

costs. In this section, we review research on how rational profit-maximizing firms respond to some richer and more realistic types of consumer preferences that have been identified in the behavioral-economics literature, and that are reviewed in other chapters of this handbook.

5.1 Loss aversion

A small literature studies market pricing when consumers exhibit expectations-based loss aversion in the sense of Kőszegi and Rabin (2006, 2007).[31] To understand the main insights in the easiest possible way, it is useful to understand the basic features of the model, and some key implications, in the context of consumer behavior. A special case of the model posits that a consumer derives utility not only from the absolute level of her consumption outcomes, but also from comparing outcomes in money and product satisfaction to her lagged rational expectations about those outcomes, with losses being more painful than equal-sized gains are pleasant. For example, if she buys a cheaper but lower-quality product than she expected, she experiences this as a loss in the product-satisfaction dimension and as a gain in the money dimension. These assumptions have three implications that are exploited repeatedly in industrial-organization applications. First, if the consumer expected to obtain a product at a lower price than the one at which it is now available, she evaluates paying the higher price as a loss, creating a "comparison effect" that makes her sensitive to changes in the price and overall less willing to buy the product. Second, because she evaluates paying a higher-than-expected price as a loss but paying a lower-than-expected price merely as a gain, she dislikes price variation to a first-order extent; and by the same argument, she dislikes variation in the quality of the product. Third, if the consumer had expected to obtain a product with a higher probability, then she evaluates not getting the product as a loss, creating an "attachment effect" that increases her willingness to pay for the product.

Implications of the comparison and attachment effects. Heidhues and Kőszegi (2008) and Spiegler (2012) show that the comparison effect can reduce or eliminate price variation in oligopolistic and monopolistic markets, respectively. In Heidhues and Kőszegi (2008), we study a Salop-type (1979) model of pricing with differentiated products, assuming that firms have uncertain cost distributions. Consumers correctly anticipate the resulting price distribution, and form rational expectations regarding their consumption outcomes before the purchase decision. The firms then optimally set prices, taking into account other firms' behavior and consumers' reference-dependent preferences. This means that firms cannot commit to their price before consumers form their reference points. Our main prediction is that prices are often identical across differentiated products—a prediction that is consistent with casual observation in such

[31] See also Grubb (2015d) for a review of this literature.

familiar markets as clothes or movies, and is documented systematically by the Competition Commission of the United Kingdom (1994), Beck (2004), McMillan (2004) and Einav and Orbach (2007). If consumers had expected to pay p^* with probability one, then they assess buying at a price greater than p^* as a loss in money and buying at a price lower than p^* merely as a gain in money, making demand kinked at p^*. Hence, for a range of cost levels p^* is the optimal price to charge, so that charging the same sticky price is often an equilibrium. More interestingly, we derive a sufficient condition under which firms with different cost distributions charge the same price with probability one in *any* equilibrium. The key step is to argue that a firm sets a deterministic price in any equilibrium; then, if the supports of firms' cost distributions are not disjoint, these deterministic prices must be the same. Suppose, toward a contradiction, that a firm's prices are—and hence consumers expect them to be—stochastic. Due to the comparison effect, demand is then more responsive at higher than at lower prices in the firm's price distribution. If the firm's costs do not vary much, it could increase profits either by decreasing high prices or by increasing low prices. We show that the tendencies for sticky and focal pricing are stronger in more concentrated industries, while in a variant of the model, Spiegler establishes that the incentive for sticky pricing is even stronger when fluctuations are in demand rather than in costs.

Subsequent work has shown that the tendency of loss aversion to generate reduced price variation depends crucially on the timing. Heidhues and Kőszegi (2014) and Rosato (2016) assume that a monopolistic retailer commits to a stochastic offer before consumers form their reference points.[32] This assumption applies if the seller can announce and commit to its selling strategy (Rosato, 2016), or can establish a reputation for playing a particular strategy (Heidhues and Kőszegi, 2014). Then, although loss-averse consumers strongly dislike uncertainty, the profit-maximizing sales strategy involves randomization even in a deterministic environment. In this result, the attachment effect plays the crucial role. Heidhues and Kőszegi (2014) assume that a monopolist with a deterministic cost sells a single product to a representative consumer with known valuation, and can announce its price distribution in advance. The optimal price distribution then consists of low and variable "sale" prices and a high and atomic "regular" price. The sale prices are chosen such that the consumer buys at these prices no matter what she had expected. Then, because the consumer expects to purchase with positive probability, the attachment effect induces her to buy also at the regular price. And as in the above papers, the regular price is sticky in part due to the consumer's comparison effect. Rosato assumes that the seller can randomize the price of the product as well as the probability that it is available. He shows that the monopolist

[32] See Karle and Peitz (2014) for a variant of Heidhues and Kőszegi's (2008) model of pricing with differentiated products in which firms commit to deterministic prices before consumers form their reference points.

optimally tempts the consumer with a bargain offer that is available in limited supply, engaging the consumer's attachment effect and thereby inducing her to buy also at the high price if the bargain is not available. Although the consumer is rational, in both of these papers the seller manipulates the consumer into a purchase that yields ex-ante expected utility below that of her outside option.[33]

First-order risk aversion and insurance. In other papers, loss aversion creates a strong (first-order) incentive for firms to shield consumers against economic risks. Herweg and Mierendorff (2013) assume that consumers are uncertain about their future demand, and show that due to consumers' dislike of price variation, the seller's optimal two-part tariff is—consistent with plenty of empirical evidence—often a flat fee rather than a measured tariff, despite such a contract inducing overconsumption. A flat fee is more likely to be optimal if marginal costs are not too high, loss aversion is strong, and there is intense variation in demand.

Hahn et al. (2018) analyze a monopolist's optimal menu when consumers are loss averse and do not know their willingness to pay in advance. To insure the consumer against fluctuations due to her range of potential willingness-to-pay realizations, the seller often offers a small number of products relative to the heterogeneity in the population.

Manipulating expectations. Since demand depends on consumers' reference points, with expectations-based reference dependence it is in firms' interest to manage consumers' expectations. This motive is present already in the models above: for instance, the reason to offer sales in Heidhues and Kőszegi (2014) and Rosato (2016) is to induce expectations of purchase. Two papers consider more direct manipulations of consumer expectations. Karle and Schumacher (2017) study a model in which a consumer is uncertain about her valuation for the firm's product. The seller can advertise to her to partially or fully reduce this uncertainty, after which the consumer forms her reference point about the purchase. Karle and Schumacher show that the advertiser may prefer partial information, so that the remaining uncertainty manipulates the consumer into making an ex-ante unfavorable purchase. The effect is similar to that in Heidhues and Kőszegi (2014) and Rosato (2016): the consumer definitely buys if her valuation is high, and due to the resulting attachment effect, she ends up buying also if her valuation happens to be lower. Karle and Peitz (2017) assume that competing firms sell differentiated products to consumers through an intermediary (e.g., Ebay). The intermediary observes the products the consumer might like, and guides the consumer by showing

[33] This last result is due to a subtle type of time inconsistency with expectations-based loss aversion. When the consumer decides to buy at a sale price, she does not take into account that this increases her ex-ante expectations to consume and spend money, lowering her expected utility.

her a set of products and their prices. The consumers form their reference points regarding the purchase, and then inspect the products to decide what to buy. Karle and Peitz establish that to increase equilibrium prices, the intermediary shows too many products. If a consumer observes a low-priced product among many higher-priced products, she assigns a low probability to eventually buying that product, so she expects to pay a higher price with high probability. As a result, she is not very sensitive to price cuts, reducing competition between sellers.

Other reference points. A few authors study the market implications of loss aversion under assumptions about reference-point determination that are not based on expectations. Carbajal and Ely (2016) posit that consumers know their types in advance and have a type-dependent reference point relative to which they evaluate outcomes. Carbajal and Ely study how the optimal menu depends on the reference-point function, and also derive properties of self-confirming reference consumption plans—where a type's consumption in equilibrium coincides with her reference point. In contrast to an individual-decisionmaking setting—where an increase in the reference point always hurts the agent—a higher self-confirming reference consumption plan can benefit both the seller and some agents. Intuitively, a higher reference point leads the seller to exclude fewer low types from the market (who, due to their higher reference point, value the product more highly), and as a result of this market expansion higher types receive higher information rents.

Zhou (2011) considers a Hotelling-type duopoly in which one of the firms is "prominent," so that its price determines the consumer's reference point in the price dimension. If the prominent firm charges a lower price than its competitor, then consumers experience buying from the competitor as a loss in money, making them really eager to avoid the competitor. If the prominent firm charges a higher price than its competitor, then consumers experience buying from the prominent firm merely as a foregone gain in money, making them less eager to avoid the prominent firm. Due to this asymmetry, the prominent firm prefers to randomize its price.

5.2 Preference for commitment

In this section, we review the literature studying market outcomes when consumers are present biased, restricting attention to situations in which consumers are sophisticated or their potential naivete is not important for the questions at hand.[34] As the drafted chapter Time Preferences for potential publication in Volume 2 of this handbook explains in detail, present-biased consumers benefit from commitment, and hence might be willing to sign contracts that restrict their choices in some way. A few papers study

[34] Since the implications of naivete regarding present bias are often similar to the implications of other types of naivete, we cover those in Sections 2 and 3.

optimal commitment contracts that result from this preference. DellaVigna and Malmendier (2004) analyze a firm's optimal two-part tariff when facing a present-biased consumer who signs an exclusive contract with the firm. The authors consider two cases. In the first case, the consumer decides whether to make an investment—such as saving or exercising—that carries an immediate cost and a future benefit. Then, the optimal tariff subsidizes marginal investment by the amount by which the consumer undervalues the benefit. In the second case, the consumer decides whether to engage in indulgence—such as eating or smoking—that generates an immediate benefit and a future cost. Then, the optimal tariff taxes the marginal indulgence by the amount by which the consumer undervalues the future cost.[35]

DellaVigna and Malmendier (2004) document that the basic features of contracts in many industries are consistent with the above predictions. For instance, consistent with the prediction that marginal investment is subsidized, health clubs often have fixed monthly fees and a marginal price of zero per visit. And consistent with the prediction that marginal indulgence is taxed, credit-card interest rates tend to be significantly above the marginal cost of credit. Note, however, that a high price for indulgence is also consistent with several models of naivete that we have discussed in Section 2: if consumers underestimate their tendency to indulge, a firm can introduce hidden prices by making indulgence expensive.

Gottlieb (2008) asks how optimal pricing to present-biased agents is modified when other firms can interact with the consumer after the initial contract is signed. He shows that for indulgence, such ex-post competition completely eliminates the firm's ability to commit the consumer's behavior. To illustrate, consider smoking. Because of her present bias, a consumer underweights the future health costs of smoking, resulting in overconsumption. She would therefore prefer to sign an exclusive contract with a firm in which she is paid a lump sum up front, and she can purchase only overpriced cigarettes from the firm in the future, restraining her overconsumption. But if she can buy cigarettes from others, the contract is completely ineffective.[36] In contrast, for investments the same concern does not arise, as other firms could not compete with the subsidized terms of the optimal commitment contract.

Nocke and Peitz (2003) study the preference for commitment in a durable-goods market with a monopolistic seller. Buying the good entails an immediate cost and a

[35] While in DellaVigna and Malmendier's setting the optimal contract achieves first-best despite asymmetric information, in general there is a tradeoff between committing a present-biased agent's future behavior and allowing her to respond to taste or income shocks. For contributions exploring optimal contracting with this tradeoff—which is unlikely to be possible in the market and therefore is outside the scope of our review—see Amador et al. (2006) and Galperti (2015).

[36] Nevertheless, in some markets for products with immediate benefits, such as credit cards, consumers do not switch easily to competitors, so (as documented by Ausubel, 1991; DellaVigna and Malmendier, 2004) high prices are still feasible.

stream of current and future benefits, so (as explained in the drafted chapter Time Preferences for potential publication in Volume 2 of this handbook) the consumer may be tempted to delay purchasing the good if she can. Since having a secondary market for the durable good facilitates procrastination, it affects outcomes in the primary market even when under classical exponential discounting it would not. As a consequence, the monopolist may want to shut down the secondary market.[37]

5.3 Markets for conspicuous consumption

The idea that individuals may want to advertise their wealth through conspicuous consumption—i.e., consumption visible to others—has been familiar to social scientists since Veblen's celebrated "Theory of the Leisure Class." In this section, we review research on how markets are affected by such a motive for consuming products. Since the products individuals use for conspicuous consumption often appear to be rather expensive, a basic question is whether and why consumers might pay supra-competitive or otherwise surprisingly high prices to signal their types. But many other questions also arise.

In most of the literature, the key assumption on the demand side is that a consumer's utility is increasing in the beliefs others hold about her wealth, income, or taste. The drafted chapter Social Preferences for potential publication in Volume 2 of this handbook reviews evidence and foundations for this particular way in which individuals like to stand out from the crowd, and discusses implications for individual decisionmaking. Pesendorfer (1995) and other research following it uses a slightly different theoretical approach, assuming that a consumer wants to impress others not because she cares directly about others' opinion, but because she wants to match with high types. As far as industrial-organizations implications are concerned, however, this motive typically leads to the same logic as the motive to stand out.[38]

Prices and product quality. In an early contribution, Bagwell and Bernheim (1996) develop a model in which competitive firms can supply different qualities of a conspicuous good, and consumers derive both an image utility and a direct benefit

[37] See also Esteban and Miyagawa (2006) and Esteban et al. (2007), who study non-linear pricing when consumers have a preference for commitment because they suffer from temptation disutility in the sense of Gul and Pesendorfer (2001).

[38] Beyond wanting to stand out, consumers may of course derive a direct, private consumption benefit from consuming the conspicuous good. A designer suit, for instance, can be used to show off, but can be valued for warmth and comfort as well. Nevertheless, for two reasons many papers assume that the conspicuous good confers no direct consumption benefit. First, this assumption captures situations in which the same benefit can be obtained from a substitute product at a much lower price. For instance, a cheap watch or a smartphone provides the same time-keeping benefits as a Rolex. Second, the assumption serves the useful theoretical purpose of isolating the implications of conspicuous consumption.

from consumption. Bagwell and Bernheim first posit a single-crossing property typically assumed in models of asymmetric information, whereby high-type consumers have higher monetary valuation for the direct benefit from increases in consumption than do low-type consumers. In this case, there cannot be an equilibrium in which high types signal by buying a product at a supra-competitive price. Intuitively, instead of spending too much on a product, a high-value consumer can buy more of a cheaper version of the product. Since she values increases in consumption more than a low type, doing so helps her signal her type, and in addition she gets more of a direct benefit from the product as well. Bagwell and Bernheim also show, however, that Veblen effects can arise if the single-crossing property is violated in a particular way. In that case, increasing consumption at the competitive price is not an effective way to signal, so a high-type consumer prefers to buy less at a supra-competitive price. And since a low type would prefer to imitate high types if the same good was cheaper, no competitor can attract consumers by lowering the price.

Mandler (2018) uses a model in which the conspicuous good has no direct consumption benefit, and higher-income consumers have a higher willingness to pay for increases in their perceived income. This framework has the notable feature that in a separating equilibrium satisfying the intuitive criterion, each consumer type's expenditure on the conspicuous good is independent of the good's price. In this sense, consumers of conspicuous goods do not care about prices. Intuitively, the role of an otherwise useless conspicuous good is purely to allow consumers to publicly burn money, and it is only how much money they burn that is used for inference. As a result, the higher is the good's price, the less is produced, so—with production costs being positive—the higher is social and consumer welfare. Piracy of conspicuous goods, therefore, can lower social welfare by lowering prices.

Kuksov and Xie (2012) ask a different question regarding how competition affects the demand for status goods. They consider a duopoly in which consumers can only purchase one unit of one product, and show that a decrease in one firm's cost can benefit *both* firms by making the competing product more exclusive and hence a better signal of high type.

While Mandler (2018) assumes linear prices and Kuksov and Xie (2012) assume unit demand, Rayo (2013) allows for both multi-unit demand and non-linear prices and asks how a monopolist optimally designs status goods—equivalently, signals of status—for consumers. Positing first that consumers derive no direct utility from the goods, Rayo shows that it is often optimal for the monopolist to pool types. Intuitively, suppose consumers within an interval of low types do not care much for increases in status, but there is a sufficiently large number of higher types that do. Then, it is optimal to pool the former consumers, selling them a product that confers low status. This allows the monopolist to extract a lot of rent from higher types by selling a product that confers high status. In an extension, Rayo assumes that consumers derive both a direct

consumption benefit and a status-signaling benefit from consumption, and higher types have higher demand for both benefits. In this case, the quality schedule is the same as in a model without status utility, but the price schedule is steeper and exhibits jumps at pools, making consumers pay a lot for high-quality-high-status goods.[39]

Friedrichsen (2017) also analyzes a model in which consumers derive both direct consumption benefits and signaling benefits from the product, but—modifying previous approaches—she assumes that the two motives are not perfectly correlated. In her model, consumers prefer to signal that they like quality. Under these assumptions, the signaling motive in general affects quality provision as well as prices, and generates a rich set of possibilities. A monopolist may offer a lower-quality product to consumers who value either quality or image, and a higher-quality product to consumers who value both quality and image. Alternatively, the monopolist may offer versions of the same product at different prices. And under competition, consumers never pay a supra-competitive price for a product, but they may buy a product with inefficiently high quality to signal taste. As a result, the market outcome may be less efficient under perfect competition than under monopoly.

Advertising and branding. Krähmer (2006) identifies a novel role for advertising in a framework related to those above. There are brand-name products and no-name products in the market, and the public initially cannot distinguish brand-name products from the rest. Brand-name producers can advertise to make their brands known, allowing consumers to signal their types to the public. Hence, advertising is aimed not at potential consumers of the product, but at the social contacts of potential consumers.

Nevertheless, it is clear that some producers of brand-name products prefer not to advertise, and even on the product itself the branding is extremely subtle and not easily identified by the average observer. Yoganarasimhan (2012) and Carbajal et al. (2016) provide different explanations for this phenomenon. Yoganarasimhan (2012) assumes that a consumer cares both about signaling her taste and about conformity—purchasing the same product as her social contacts.[40] An inconspicuous branded product can be identified only by sophisticated consumers, and hence is particularly helpful for signaling taste. A conspicuously branded product, in contrast, helps a consumer conform to a greater share of possible contacts. Carbajal et al. (2016) assume that a consumer wishes to signal not only her wealth, but also her social connectedness, and inconspicuously branded products are recognized only in close social interaction. In contrast to a flashy car, for instance, a painting in one's living room is only seen by one's guests. Then a

[39] See also Mazali and Rodrigues-Neto (2013) for a related analysis of product offerings where pooling is driven by fixed costs of producing a brand rather than price discrimination, and the implications for tax policy.

[40] See Amaldoss and Jain (2005) for a comparison of the implications of the desire for uniqueness versus conformity.

wealthy and socially well-connected individual can differentiate herself from wealthy but socially unconnected others by purchasing inconspicuous status goods: although the latter consumer could afford such products as well, they would not make sense for her to buy, as not many would see it. Because inconspicuous brands therefore confer the highest status, they are—consistent with casual observation—more expensive than flashy alternatives.

Dynamics. A few papers have explored the dynamics of markets for conspicuous consumption. Pesendorfer (1995) provides an explanation for fashion cycles in a model where a monopolist can at any time pay a fixed cost to introduce a new version of the product, but cannot commit to then refrain from selling it in future periods. The monopolist periodically designs a new version of the product, and sells it expensively to allow consumers to signal their types. This signaling value becomes diluted as the product spreads in the population, setting the stage for a new fashion. Kuksov and Wang (2013) provide an explanation for why fashions appear to be partly random: to make it difficult for low-type consumers, who cannot quickly learn about new fashions, from purchasing the hit product of the day. Relatedly, Rao and Schaefer (2013) show that the drop in a monopolist's price over time is greater for status goods than for classical goods, and they discuss various commitment strategies, such as product dating or product changes, whereby a firm can limit this effect for early adopters.

Amaldoss and Jain (2008) consider the interaction between leaders who are discerning enough to purchase the status good early and followers who only get the chance to purchase late. Consumer preferences are not directly driven by signaling type; instead, the authors assume that leaders' willingness to pay is decreasing in the number of followers they expect to purchase, and followers' willingness to pay is increasing in the number of leaders who have purchased. A firm then faces a commitment problem: once leaders have purchased, the firm wants to sell to followers, but the expectation that it will do so lowers leaders' demand. If the leader market is important, the firm charges a high price early and reduces demand from leaders to make it unattractive to sell to followers. In contrast, if the follower market is important, then the firm lowers prices to leaders so that they purchase despite the expectation that followers will also buy. Because the firm may want to commit to selling little to followers, it might choose a limited-edition product, or an overly fancy product with a high marginal cost.

6. BEHAVIORAL MANAGERS AND FIRMS

Much of the behavioral-industrial-organization literature—and hence much of our review—discusses the implications of psychologically based models of demand rather than psychologically based models of supply. In part, this tendency reflects a conscious judgment by researchers about what is important: profit-oriented firms have incentives to safeguard against profit-decreasing psychological tendencies by employees, and

therefore many phenomena are likely to be less prevalent in the behavior of firms. Nevertheless, there are at least two major reasons to expect psychological phenomena on the supply side as well. First, a small owner-managed firm does not have the capacity to put in elaborate safeguards against profit-decreasing tendencies, so whatever psychological phenomena the owner exhibits as a consumer surely manifest themselves in the firm's behavior.[41] Second, some psychological tendencies might be profitable for a firm—even if they are utility-decreasing for the employee—or might predispose a person toward starting a business or corporate career, so that selection effects do not necessarily eliminate all psychological tendencies. Indeed, while it is much smaller than the literature on behavioral consumers, there is a growing literature on behavioral tendencies in firms, which we discuss in this section. Unfortunately, the literature has so far studied only limited aspects of managerial behavior. Given the many business decisions a typical firm undertakes, it seems that there is a lot of room to use behavioral models to improve our understanding of firms' decisions.

6.1 Firm mistakes

We first discuss types and sources of mistakes—that is, suboptimal decisions—by key decisionmakers in firms, beginning with small owner-managed firms and continuing with larger firms. In line with the selection criteria for this survey, we do not discuss work that documents mistakes by firms without considering the underlying psychological sources.[42] Furthermore, we do not discuss what market outcomes obtain if firms follow simple learning rules such as reinforcement learning, because we think of these as descriptive rather than psychologically rich models of firm behavior.[43]

[41] In a well-put observation emphasizing that we cannot expect entrepreneurs to just maximize profits, Axinn (1983) noted: "an economist will forcefully express the view that the only meaningful goal of the rational business executive is the maximization of his own profits . . . that is not going to ring true to anyone who has . . . had to put his son-in-law in a business."

[42] In the classical industrial-organization literature, profit maximization is typically *assumed*, not documented. One reason is that to test the profit-maximization hypothesis, a researcher needs to know the firm's cost *function*, which is typically not available. Nevertheless, a number of papers document failure to maximize profits. For example, Cho and Rust (2010) demonstrates through a field experiment that rental-car companies could increase profits by holding on to their cars longer while giving consumers a discount for renting older cars; Hanna et al. (2014) document that seaweed producers fail to optimize with respect to the initial seaweed pod size; DellaVigna and Gentzkow (2017) report that US retail chains give up about 7% in profits by often charging the same (uniform) price across outlets despite differences in the demand and competitiveness of local markets; Covert (2015) provides evidence that firms in the US shale industry respond too little to new information—especially if the data originates from competitors—and fail to exploit profitable experimentation; Bloom et al. (2013) show in a field experiment that the introduction of better management practices significantly improves firm performance in the Indian textile industry; and more generally, the literature on "management as technology" finds that some firms consistently use inferior management practices (Bloom and van Reenen, 2007).

[43] See Armstrong and Huck (2010) for an in-depth discussion of these papers.

Small firms. Available evidence indicates that starting a small business is unlikely to be financially rewarding: the majority of businesses fail quickly (see, for instance, Artinger and Powell, 2016, and the references therein), and the majority of entrepreneurs would be financially better off with a salaried job (Hamilton, 2000; Moskowitz and Vissing-Jørgensen, 2002). Although other phenomena surely contribute to this pattern and the evidence on the relative importance of different factors is far from conclusive, many researchers suggest that owners' overconfident beliefs regarding success play a crucial role.[44] For instance, Manove and Padilla (1999) observe that (all else equal) an optimist is more likely to found a firm than a pessimist, so we would expect overconfidence to be overrepresented among entrepreneurs. Indeed, people starting small businesses think their business is far more likely to succeed than a typical similar business (Cooper et al., 1998), and entrepreneurs of startups have unrealistically optimistic beliefs regarding future growth (Landier and Thesmar, 2009).

Motivated by the above considerations, two papers study the implications of overconfidence for debt financing. Manove and Padilla (1999) is an early paper that develops a model of project selection by entrepreneurs who tend to overestimate the profits they will make. In the model, entrepreneurs receive an informative signal regarding future profits, and while a realistic entrepreneur interprets the signal correctly, an overconfident entrepreneur always interprets it as being good. Entrepreneurs then choose whether to make a small or a large investment, and go to a competitive market to finance their investments. The authors look for perfect Bayesian equilibria that satisfy the intuitive criterion, adjusted for the fact that overoptimistic entrepreneurs misinterpret their signal. Comparing the equilibrium outcome with the second-best, they find that market financing is not conservative enough. To see this, note that an overoptimistic entrepreneur may want to engage in a high investment despite having a bad project. The market interest rate reflects the true population risk and ensures that the project generates enough in expected profits to be worthwhile, but it does not correct for the opportunity cost of the investment: the overoptimistic entrepreneur could have undertaken a low-investment project, which would have generated higher social surplus in expectation. Furthermore, collateral requirements (or unlimited liability) do not (sufficiently) deter overoptimistic entrepreneurs who believe their project is unlikely to fail. In contrast, bankruptcy requirements and other limited-liability mechanisms will raise equilibrium interest rates and deter investments based on unrealistic optimism.

Landier and Thesmar (2009) assume that an overconfident entrepreneur overestimates the probability of success when starting her business, and ignores the fact that

[44] Alternative explanations include preferences regarding the gamble involved in starting a small business, as well as the non-pecuniary benefits of self-employment. Hall and Woodward (2010) estimate, for instance, that because entrepreneurs backed by venture capital have a small probability of making a really large payoff, individuals who have high initial assets or low coefficients of relative risk aversion may rationally choose to become entrepreneurs.

low interim cash flow is a bad sign about profitability. As a result, overconfident entrepreneurs sign short-term debt contracts. Intuitively, an overconfident entrepreneur believes that a low interim cash flow is unlikely and hence short-term debt is a good deal, and investors in turn value the potential to liquidate bad projects at the interim stage. Because realistic entrepreneurs are willing to liquidate bad project by themselves, in equilibrium they choose long-term debt, which provides better hedging benefits. Landier and Thesmar document that consistent with this prediction, more optimistic entrepreneurs are more likely to use short-term debt for financing.

Based on the distinction suggested by Moore and Healy (2008), Astebro et al. (2014) emphasize that different types of overconfidence should have different implications for markets. If overconfidence amounts to overestimation of one's own ability or performance, then—similarly to the overconfidence assumed by Manove and Padilla (1999) above—it may drive individuals to entrepreneurship if ability has a relatively higher return for entrepreneurs. If overconfidence amounts to overplacement—according to which individuals overvalue their skills relative to others—then it may induce overentry into competitive markets in particular, with too many entrepreneurs thinking that they can outperform others. If, on the other hand, overconfidence amounts to overprecision—according to which individuals have excessively narrow confidence intervals around their estimates—then overconfident entrepreneurs will often undervalue the benefit of exploration relative to exploitation (Herz et al., 2014), and will be likely to ignore feedback and stick with their chosen paths. But in contrast to other forms of overconfidence, one would not expect overprecision to bias individuals towards attempting entrepreneurship.

Studying a different mistake, Goldfarb and Xiao (2017) argue that inexperienced restaurant owners in Texas make a relatively small, but predictable error in their exit decisions: they fail to condition on the transitory nature of weather shocks, leading them to exit too early after bad weather shocks and too late after good weather shocks. Experienced owners, in contrast, properly condition on weather shocks. Adapting the sparsity model of Gabaix (2011) (see the chapter Behavioral Inattention that is in preparation for Volume 2 of this handbook), Goldfarb and Xiao (2017) develop and estimate a structural model in which attention costs can prevent owners from taking weather shocks into account, and find evidence that between 83% and 91% of owners do not pay attention to weather at all, with experience leading to a sharp drop in the estimated attention costs.

Large firms. In large, publicly traded firms, harmful individual behavioral tendencies may be mitigated through cooperate governance structures and the selection of capable leaders, so one may conjecture that some phenomena—e.g., procrastination—will be less relevant than in small firms. But any sweeping claim that behavioral tendencies cannot be important for managerial decisions is theoretically and empirically

misguided. From a theoretical perspective, Simon (1955) has argued long ago that given how difficult a large firm's optimization problem is, profit maximization is an implausible assumption for the "administrative man." Rather, firms' managers optimize "locally" and try to achieve "satisficing" outcomes.[45] In addition, there may be behavioral tendencies that a firm profits from overall, so that the firm might seek out—and potentially exploit—candidates with these traits. As a notable example, de la Rosa (2011) shows that a firm benefits from an overconfident employee, because it can reduce the expected wage through a performance contract and possibly also implement high effort more efficiently. And the corporate promotion process may also disproportionately favor some behavioral traits. Again, overconfidence can help a person live through the cutthroat competition involved in becoming a top manager. Consistent with these arguments, a sizable empirical literature, to which we now turn, documents some psychological sources of mistakes by managers of large corporations. We suspect that the findings have some broader implications as well, for instance for the dramatic and persistent within-industry productivity differences among firms with access to the same resources (Syverson, 2011), or the effects of CEO characteristics on firm performance (Bertrand and Schoar, 2003), but have not seen these possibilities explored in the literature.

An old hypothesis in industrial organization and finance is that unprofitable mergers are partly driven by empire-building preferences, or managers' hubris in their ability to evaluate (Roll, 1986) or run an acquired firm. Malmendier and Tate (2005, 2008) reinvigorated this hypothesis by documenting that managerial characteristics—especially overconfidence—influence firm decisionmaking in major ways. We refer the reader to the chapter Behavioral Corporate Finance of this handbook for details on this issue, as well as how it interacts with optimal corporate governance. Here, we discuss implications of manager traits for pure industrial-organization questions.

The hubris hypothesis for takeovers is closely related to the winner's curse in auctions, which Capen et al. (1971) and later Thaler (1988) proposed as an explanation for firms' consistently low returns on winning offshore oil leases. As Capen et al. (1971) put it well before the ascent of auction theory: "There is a somewhat subtle interaction between competition and property evaluation, and this phenomenon—this culprit—works quietly within and without the specific lease sale environment. We would venture that many times when one purchases property it is because someone else has already looked at it and said, 'Nix.' The sober man must consider, 'Was he right? Or am I right?' " This failure of strategic reasoning—ignoring the information contained in others' actions—is well-documented in experiments and central in behavioral game theory, and is discussed in depth in the chapter Errors in Strategic Reasoning that is in preparation for Volume 2 of this handbook. Yet the striking feature about the original

[45] See also Simon (1959) and the discussion of boundedly rational firm behavior in Ellison (2006) and Armstrong and Huck (2010).

oil-lease context is that it is a high-stakes business decision in which bidding firms had ample feedback. Is it possible that oil firms invested large sums in costly exploration and estimation of oil reserves, only to slip due to a failure in basic strategic reasoning? It seems so: Hendricks et al. (1987) find evidence that by shading bids in all auctions by a constant fraction, roughly three-quarters of the firms could indeed have increased their profits—suggesting that they may indeed have fallen for the winner's curse. Consistent with strategic mistakes in bidding behavior, Hendricks and Porter (1988) find that firms who do not own neighboring tracts—which leads to an informational disadvantage at predicting a tract's value—lose money on average. While the estimated losses are insignificantly different from zero, if non-neighboring firms have some private information regarding the tract's value, due to information rents they should earn positive profits. And even absent such private information, an optimizing firm must earn positive profits conditional on winning in order to recoup its bidding cost.

More recently, Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) structurally estimate managers' strategic ability in technology adoption and entry decisions by applying the cognitive-hierarchy approach of Camerer et al. (2004) (discussed in detail in the chapter Errors in Strategic Reasoning that is in preparation for Volume 2 of this handbook). They define a level-zero player as someone who conditions her choice on publicly available information, but not on rivals' anticipated decisions—that is, she acts as a potential monopolist would. Higher-level players, in turn, behave as if all other managers are of lower levels, with those levels drawn from a truncated Poisson distribution. Because higher types can make better predictions about rivals' play, Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) interpret the estimated cognitive hierarchy level of a manager as her strategic ability.

Goldfarb and Yang (2009) investigate technology adoption by internet service providers, while Goldfarb and Xiao (2011) investigate entry decisions into local US telecommunication markets following the deregulatory Telecommunications Act of 1996. Holding other market characteristics constant, one key reduced-form finding in Goldfarb and Xiao (2011) is that better-educated managers—those with degrees from very good undergraduate institutions, or degrees in business or economics—tend to enter markets with fewer competitors, which suggests that they are better at predicting competitors' behavior. In the structural analysis, Goldfarb and Xiao's (2011) key idea is to exploit the variance firms display in entering markets. In particular, note that type-0 managers (who are acting as monopolists) have a high probability of entering a given market, so that type-1 managers (who are optimizing against type-0 managers) have a low probability of entering, and type-2 managers (who are optimizing against a mix of type-0 and type-1 managers) have an intermediate probability of entering. This implies that an intermediate overall probability of entry is consistent with a combination of type-0 and type-1 managers as well as a predominance of type-2 managers, but the former case exhibits higher firm-specific variation in entry probability. The authors

estimate that better-educated managers have higher strategic ability. Furthermore, the estimated level of cognitive ability is higher in 2002 than right after the first wave of entry in 1998, suggesting that low-ability managers were more likely to fail in the intervening shake-out. And reassuringly, both Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) find that a manager's estimated cognitive ability is positively correlated with revenues and the probability of staying in business out of sample.

Hortaçsu and Puller (2008) analyze the bidding behavior of electricity firms in the Texas "balancing market," where suppliers trade between each other to meet prior contractual obligations. Large firms best-respond to other firms' behavior, but small firms submit excessively steep bid functions, insufficiently adjusting their production quantities to market circumstances. While there is some learning by small firms, the learning rate is relatively low. The authors argue that the cost of setting up a sophisticated bidding unit is the main reason for small firms' underparticipation in the balancing market. Furthermore, they estimate that the differential bidding behavior by firms significantly reduces productive efficiency in the marketplace—suggesting that taking firms' mistakes into account can be an important consideration in designing electricity (and likely other) markets.

Taking a more micro-founded approach, Hortaçsu et al. (2017) adapt the cognitive hierarchy model to bidding in the Texas electricity market. They assume that a level-zero firm simply submits a perfectly inelastic bid function at its contract position, and a higher-level firm best-responds to a truncated Poisson distribution over lower-level types, taking into account its believed market power. Because higher-level firms believe their rivals to be higher-level as well, they face a more elastic residual demand curve, leading them to bid more competitively. Hortaçsu et al. (2017) illustrate that increasing sophistication—either exogenously or through a merger—leads to efficiency gains because higher-level firms therefore bid closer to marginal cost.[46]

Massey and Thaler (2013) investigate teams' behavior in the NFL draft, a labor market with two unique features that make it ideal for identifying mistakes. First, the quality of hiring decisions can be measured unusually well based on players' subsequent performance. Second, due to the salary cap—a cap on how much the team can spend on players' wages in total—an owner's problem can be thought of as using a "given budget to buy the best performance." Massey and Thaler argue that a number of psychological

[46] Doraszelski et al. (2018) investigate bidding behavior, and especially learning, in the newly introduced market for frequency response within the UK electricity system. They argue that as initial demand and rivals' initial bidding behavior are difficult to predict, early on there is prone to be considerable strategic uncertainty. Indeed, at the beginning there are frequent and sizable adjustments in bids, and bidding behavior is not in line with equilibrium predictions. But in contrast to the slow learning documented by Hortaçsu and Puller (2008) for the Texas balancing market, already after one and a half years play starts to converge towards complete-information Nash-equilibrium predictions, and reaches these within four years after the market is opened.

forces can lead NFL teams to overvalue early draft picks relative to later ones.[47] These forces include failing to account for reversion to the mean—that an exceptionally good performance is often followed by more mediocre performance—and the winner's curse, as well as overconfidence in being able to identify top performers. Indeed, while early draft picks perform better than later draft picks, their higher performance does not warrant the steeply higher wages they are paid. In what Massey and Thaler term the loser's curse, the first draft pick (which goes to the previous season's worst performer) generates lower surplus than any second-round pick—contradicting the rationality of the market.[48]

6.2 Behavioral contracting and the theory of the firm

In their classic textbook on industrial organization, Scherer and Ross (1990) discuss the case of the American steel cartel. According to Judge E. H. Gary (chairman of US Steel's board of directors), social interaction during the dinners he hosted helped establish trust (or "respect and affectionate regard") among the steel industry's leaders that lead to cartel agreements being "more binding . . . than any written or verbal contract." The quote suggests that social preferences—partly formed through social interactions—are crucial for understanding how firms interact. While the old implicit challenge to model social motivations in collusive settings has not been met yet, a series of papers investigates how negotiations between firms—or their managers—are influenced by psychological tendencies such as (negative) reciprocity.

Conceptually, it is useful to relate the novel behavioral approaches to the classical property rights approach to the theory of the firm (Grossman and Hart, 1986; Hart and Moore, 1990). The property rights approach assumes that while it is possible to specify asset ownership and make transfers ex ante, due to the complexity of the environment one cannot completely specify outcomes for all contingencies ex post. The resulting incomplete contracts give parties an incentive to renegotiate ex post. Crucially, ex-post renegotiation is assumed to be efficient. Thus, the only role of asset ownership in these models is to influence non-contractible interim investment incentives. While the literature has generated a number of very important insights, it is hard to imagine that the only purpose of asset ownership is to motivate interim investment. Furthermore, casual observation suggests that renegotiation of contracts is neither ubiquitous nor friction-

[47] New players, so-called rookies, enter the NFL through the "draft." Unless traded, a drafted player can only play in the NFL for the team that drafted him. Teams select players in a predetermined order, with the team owning the first draft pick selecting first, followed by the team owning the second draft pick, etc. One round of the draft ends once all teams have had their turns. Crucially, the draft picks can be and are traded among teams. Massey and Thaler assume that if two teams trade, say, the first pick in exchange for the 12th and 34th picks, then these are valued (approximately) equally by the teams.

[48] See also Romer (2006), which documents that NFL teams are too reluctant to go for a first down over kicking on fourth down. The reason for this behavior is unclear.

less. Incorporating psychological tendencies such as negative reciprocity, loss aversion, and self-serving biases has given rise to models in which renegotiation is not necessarily efficient. In such circumstances, initial contracts are designed in part to reduce negotiation inefficiencies ex post, and as we will see it is sometimes better not to specify any contract ex ante. The renegotiation inefficiencies, thus, reintroduce some of the considerations that were prominent in the earlier transaction-cost approaches to the theory of the firm (Coase, 1937; Simon, 1951; Williamson, 1971; Klein et al., 1978). But the novel microfoundations lead to new predictions that promise to enrich our understanding of how production is organized, both within firms and in long-term relationships between firms.

Renegotiation and shading. Hart and Moore (2008) initiated the behavioral approach by introducing a number of psychological assumptions into a contracting problem between a buyer and a seller. In their model, a contract written at date 0 determines the parties' entitlements at date 1. Falling short of one's entitlement leads to a loss, which can be reduced by costlessly shading one's performance to the detriment of the other party. Importantly, the parties have a self-serving view of their entitlements: they feel entitled to the best contractually feasible outcome.

The above framework generates a central tradeoff: a rigid contract guarantees that each party receives what she feels entitled to and thereby prevents shading; but a rigid contract cannot be contingent on information that arrives later and is therefore ex-post inefficient. Hart and Moore (2008) predict that parties are more likely to restrict those aspects of the contract—such as the price—over which there is a strong conflict of interest, while other variables may be specified to a lesser extent. They also predict that to reduce the amount of shading, the party who has a stronger preference regarding the design of the product should be allowed to specify it ex post. In the extreme case in which one party is almost indifferent between different specifications while the other cares a lot, a fixed-wage "employment relationship" results: the employee—who cares little about the exact task she has to perform—does the task decided upon by the employer—who feels strongly about the exact product design.

Hart and Holmstrom (2010) build a related model to identify a novel tradeoff regarding the optimal scope of the firm. Each of two units has a "boss" who implements a binary decision of whether or not to "coordinate." Joint profits of the two units are maximized if both units coordinate. Each unit's boss, however, also receives a non-transferable private benefit from the chosen activity, and so total surplus—the sum of profits and private benefits—may or may not call for coordination. The authors assume that profit sharing is impossible, and that renegotiation is not practicable due to the threat of shading. They compare two decision structures: either each unit's boss can make the decision of whether to coordinate (non-integration), or there is an outside manager with the aim of maximizing total profits who makes the decisions (single-firm

integration). On the one hand, if the benefits from coordination are unevenly divided, then non-integration leads to too little coordination. On the other hand, as long as coordination leads to a reduction of private benefits, single-firm integration leads to too much coordination. And in either case, when a unit manager does not like what happens, she may shade, worsening the inefficiency.

Hart (2009) introduces uncertainty regarding the cost and benefit of trading into a model in which ex-post renegotiation is plagued by shading behavior. In his model, a buyer and a seller can agree ex ante to a fixed (ex-post) price.[49] When the state of the world is realized, each party decides whether to trade at the pre-specified price or to hold up her trading partner by insisting on a price adjustment. Because a forced renegotiation results in shading behavior, a party chooses hold-up only if the renegotiated price is sufficiently better than the one specified in the contract. The model, hence, predicts that renegotiation and shading are more likely to occur in volatile settings. Furthermore, the renegotiated price depends on parties' outside options, which are partly determined through asset ownership. Because asset ownership increases the outside option when the state of the world is good, Hart predicts that one should assign asset ownership to the party with a more state-sensitive valuation from trading. Doing so gives the party higher bargaining power exactly when she gains more from trading, making it less profitable to hold her up and reducing the occurrence of inefficient shading behavior. This prediction contrasts with that from the classic property-rights approach, where asset ownership is assigned solely to increase non-contractible interim investments.

Renegotiation under loss aversion. Herweg and Schmidt (2015) follow Hart and Moore in assuming that a contract acts as a reference point that parties dislike falling short of, but they posit that the source of this dislike is loss aversion rather than a biased view of entitlements. Consider a buyer and seller who negotiate over a good to be delivered at a later point in time, and who are both loss-averse over two dimensions of utility: for the buyer a money and a product-satisfaction dimension, and for the seller a money and an effort-provision dimension. In this setting, a contract that specifies the price and the product to be traded makes—akin to the well-known endowment effect—parties reluctant to switch to a different trade. This implies that if the ex-post optimal terms are close to the specified ones, parties do not renegotiate; and even otherwise, they only partially adjust the contract terms. Hence, in sharp contrast to Hart and Moore (2008), here it is a specific contract that leads to a renegotiation inefficiency. To avoid setting a utility-decreasing reference point, therefore, it may be better for parties not to write a contract. Similarly, even if writing a contract is optimal, parties may agree on a "compromise" contract that is never efficient ex post but that limits the ex-post

[49] He also considers an extension in which parties can specify a price range, which through the same mechanism as in Hart and Moore (2008) comes at the cost of inducing shading.

renegotiation inefficiency. Herweg and Schmidt (2015) also compare an at-will employment contract—in which the buyer can order a specification (as in Simon, 1951), but the seller is free to walk away—to a fixed performance contract. The optimal contract is determined by the scope for inefficient abuse generated by an employment contract and the renegotiation costs generated by a specific contract.

Herweg et al. (2018) develop a closely related model in which a buyer who is loss-averse in the sense of Kőszegi and Rabin (2006) trades with a profit-maximizing seller in an incomplete-contracting environment. There is a single good to be traded, and ex ante there are three possible specifications. Depending on the state of the world, each specification turns out to be either useless, to generate a low value for the buyer, or to generate a high value for the buyer. The seller costs depend on the value generated to the buyer and are such that it is always efficient to trade the low-value good. An ex-ante contract specifies a price and assigns the right to select a specification to one of the parties. Absent loss aversion, ex-post bargaining always ensures that the materially efficient specification is traded, so the contract is irrelevant. But loss aversion can render some contracts inefficient by creating an expectation to trade inefficiently.[50] In contrast, a seller employment contract always remains optimal. Because the buyer strongly dislikes the worthless (but cheap-to-produce) specification, renegotiation occurs and the materially efficient good is traded ex post. But then for any state of the world the buyer can foresee the ultimate payment as well as the valuation from the good she will consume, so this contract also induces no loss.

6.3 Firm behavior in markets: motives beyond profit maximization

Market measures beyond profits. There is some classic work in industrial organization on firms' objectives. Much of the literature derives deviations from the profit-maximization motive from the contracting problem between an owner who is solely interested in maximizing her wealth and a manager she hires to run her firm.[51] In these problems, it is optimal to make the manager's compensation contingent on outcomes other than output, leading the manager to pursue those broader objectives. To optimally induce unobservable effort from a risk-averse manager in the presence

[50] To see this, consider an employment contract in which the buyer has the right to choose her preferred specification. Then absent renegotiation, she always selects the inefficient high-value specification. To reach material efficiency, the seller can offer a price reduction in exchange for trading the efficient low-value specification. If the buyer expected to go along, she would expect to always pay the lower price, and would therefore feel a loss when paying the contractually agreed price. To reduce this loss, she would be willing to accept a smaller price cut. Hence, the buyer may prefer not to expect to renegotiate, and if she is sufficiently loss averse this can be credible.

[51] An exception is Baumol (1958), who hypothesized based on casual observation that the typical American corporation maximizes sales subject to profits reaching an acceptable level. He argues that this alternative model can explain otherwise puzzling observations, such as that increases in overhead costs are passed on to prices.

of industry-wide shocks, for example, Holmström (1979, 1982) shows that an element of relative performance pay is optimal. Furthermore, in oligopolistic industries in which a manager's incentive contract is or can be made known to rivals, an owner often strategically delegates: in an attempt to influence rivals' behavior, she rewards the manager based not only on profits, but also on output (Vickers, 1985) or market share (Fershtman and Judd, 1987). This enables an owner to behave as a Stackelberg leader: he can—through writing the appropriate goals into the contract—induce her manager to choose the Stackelberg-leader action.[52]

While research on strategic delegation focuses on contracting, one can straightforwardly extend the logic to the selection of managers whose preferences are known to rivals.[53] For example, strong status concerns with respect to rivals will lead a manager to act more aggressively, potentially making her a good hire. To see this intuitively, suppose that a manager who cares about relative profits leads a firm in a symmetric Cournot game. Since a slight increase in output has a second-order effect on the profits of her firm but a first-order negative effect on the profits of rival firms, she will act more aggressively. This induces rivals to produce less, and thereby increases her own firm's profits. But owners do not always want to select aggressive types. With standard differentiated-products price competition, an owner would like to commit to non-aggressive behavior and, thus, avoid selecting an aggressive type. Similarly, an owner who hopes to profit from collusion will often want to avoid a manager driven by relative profit concerns.

Shifting the research focus from exploring reasons behind firm motives to exploring the implications of specific motives, Cabral (2018) assumes that a firm likes to be number one in terms of market share.[54] Two firms interact over an infinite horizon, with consumers subject to taste shocks infrequently and randomly reconsidering from which firm

[52] Depending on the nature of the market game, this may increase or decrease competition. In a regular Cournot market, owners will want to reward higher output or sales to increase profits. In contrast, in standard differentiated-products price-competition models, the strategic complementarity implies that owners want to give incentives to set higher prices (Tirole, 1988). Taking a more general, game-theoretic approach, Heifetz et al. (2007) show that for almost every game, a player materially benefits from committing to maximizing something other than her true preferences.

[53] This motive is conceptually identical to that in Schelling's (1960, pp. 142–143) famous observation: "The use of thugs or sadists for the collection of extortion or the guarding of prisoners, or the conspicuous delegation of authority to a military commander of known motivation, exemplifies a common means of making credible a response pattern that the original source of decision might have been thought to shrink from or to find profitless, once the threat had failed."

[54] While it is not his main research question, Cabral does observe that the motive to be number one might have a strategic advantage. To see this, suppose firms compete for market share in period 1 and then compete for customers who incur a switching cost in period 2, and consider the firm that likes being number one. If this firm does not attract a sufficient number of consumers in period 1, it is bound to price aggressively in the period 2 to become number one. The firm's rival, therefore, has an incentive to price less aggressively in period 1 to avoid a future price war. This leads to higher equilibrium profits for the firm.

to purchase. Firms use Markov pricing strategies that condition on their current market shares, and in the simplest case both firms have a preference to be number one. Then, firms tend to price aggressively when market shares are close to equal and hence market leadership is up for grabs, but not when market shares are asymmetric. To decrease the chance of price wars, therefore, the firms have a mutual interest to allow a market leader to increase its market-share advantage, so that market shares can stay asymmetric for long periods. Unlike in models with increasing returns, however, the industry does not tend to permanently tip in one direction, as fortunes can reverse after a price war.

Corporate social responsibility. There is a wide-spread belief—especially outside economics—that firms should engage in socially desirable acts over and above those implied by laws and contracts. The economics literature on such "corporate social responsibility" (CSR), even broadly interpreted, is relatively small.[55]

Bénabou and Tirole (2010) contrast three possible interpretations of the term: an undertaking motivated largely by profit maximization, albeit over a longer horizon than in other business decisions; engagement in pro-social behavior on behalf of stakeholders; and insider-initiated corporate philanthropy. Bénabou and Tirole argue that despite skepticism from economists, some of these types of CSR might make sense, although they also highlight that especially the third type is highly controversial. The first type of CSR may be needed to counteract the potentially excessive short-term focus induced by corporate governance inefficiencies; and the second type can be called for if transaction costs are lower when firms rather than stakeholders engage in a pro-social activity.

Predating the modern CSR debate, Kahneman et al. (1986) argue that consumers feel entitled to a fair share of the benefits of trade. To investigate the implications for industrial organization, Rotemberg (2005) develops a model in which all firms are selfish, but consumers believe that some altruistic firms may exist. A consumer is willing to buy from a firm as long as she cannot reject that the firm is altruistic, i.e., whenever the firm's price falls below a reference price determined by her personal belief regarding the firm's cost. Then, the firm sets either the monopoly price (if this does not antagonize any consumer), an interior price below the monopoly price, or a price just below the one that starts antagonizing consumers. In the last case, the firm faces a kinked demand curve. In the dynamic extension, consumers reevaluate a firm only if it changes its price, making firms reluctant to change prices as costs change.[56]

[55] There is a vast management literature on the topic that we do not review.

[56] In Rotemberg (2011), a consumer believes that some firms are altruistic and naive—in that they ignore customer anger when determining prices—and are willing to punish a firm if they perceive it to be insufficiently altruistic. Rotemberg identifies conditions under which even selfish firms mimic naive altruistic ones to avoid antagonizing customers, and shows how this may play out in pricing decisions under various assumptions. Rotemberg (2010) analyzes a closely related model of dynamic price adjustments with customer regret and argues that this can better explain evidence relating price adjustments and

Baron (2009) investigates conditions under which firms may incur costs to reduce an undesirable externality. He develops a duopoly model in which a selfish firm competes with a socially responsible firm, which dislikes producing the negative externality. Citizen consumers are heterogenous regarding how much they are willing to pay for a firm's "corporate social performance" when buying its product. In equilibrium, the selfish firm engages in zero corporate social performance while the responsible firm engages in a positive amount both because it cares intrinsically and because this lowers price competition through increasing vertical product differentiation. Baron also considers the case in which citizen consumers can contribute to an activist who aims to maximize corporate social performance. The activist can demand additional corporate social performance from at most one of the firms by (credibly) threatening to carry out a costly campaign if and only if her demands are not met. Interestingly, in the baseline model the activist wants to target the responsible firm. Intuitively, the selfish firm (but not the responsible firm) is difficult to motivate because engaging in corporate social performance carries with it the indirect cost of lowering product differentiation and thereby intensifying price competition. When consumers care about the baseline corporate social performance—and are therefore less willing to follow an activist in punishing a morally-inclined firm—then it can be optimal to target the selfish firm.

7. POLICY ISSUES AND IMPLICATIONS

We discuss a few general principles that emerge from work in behavioral industrial organization regarding the regulation of markets. Some of these principles have already found their way into actual policies or policy proposals in specific market contexts (see Nuñez, 2017, for examples). We expect an increasing number of future competition and consumer-protection cases in which behavioral-economics insights are central, as well as a flood of cases in which behavioral-economics insights help qualify or inform a more classical analysis.[57] But discussing regulatory tradeoffs in specific markets is beyond the scope of this review, so we restrict attention to general principles.

7.1 The inadequacy of market solutions to consumer mistakes

Competition. A famous quote from a former chairman of the Federal Trade Commission is that "robust competition is the best single means for protecting consumer interests" (Muris, 2002 cited in Spiegler, 2015). This claim presumably reflects the idea

inflation. Customer regret following unanticipated price changes can also explain why firms announce price changes in advance.

[57] For instance, the German antitrust authority is considering a case against Facebook arguing that the company's collection and handling of private data amounts to abuse of dominance. This claim is based in part on the idea that consumers do not understand what private data they are giving to Facebook, and how Facebook will use it (Bundeskartellamt, 2017).

that a competitive market both supplies a given product at cheaper prices, and supplies more efficient products, than a less competitive market. While the claim might be broadly correct with fully rational consumers, the models we have discussed imply that when consumers make mistakes, competition and other market-based solutions often do not help consumers.

Among the classical advantages of competition, the one that generalizes best to models of behavioral industrial organization is perhaps the low level of prices. In particular, the broad qualitative conclusion of many models in Section 2 is that competition reduces prices. Even when it comes to prices, however, consumer mistakes in comparing quality (e.g., Spiegler, 2006b; Gamp and Krähmer, 2017) or prices (e.g., Spiegler, 2006a) can soften price competition in multiple ways, dampening the price-reducing power of competition and even encouraging firms facing fiercer competition to obfuscate more (e.g., Carlin, 2009).

When it comes to providing efficient products, the literature we have reviewed indicates that competition is unlikely to provide any help on average. Competition may decrease prices without any effect on exploitation distortions (as in the basic cases we have discussed in Section 2.4), it may increase exploitation distortions (e.g., Gamp and Krähmer, 2017), and it may increase or decrease the incentive to educate consumers about product quality (e.g., Heidhues et al., 2017).

Advice from intermediaries. But markets can help consumers in ways beyond the direct reduction of prices or improvement in quality. A popular narrative is that even if consumers cannot navigate some complex market environment by themselves and firms choose not to educate them, consumers can turn to expert advisors for help. Such information intermediaries could help consumers make better decisions, albeit for a fee. Price comparison websites, which allow consumers to find cheap flight tickets or hotel rooms, may be an example of such a helpful intermediary (Kamenica et al., 2011). At the same time, there are reasons to doubt that intermediaries can eliminate consumer mistakes in markets.

An obvious problem is firms' reaction to the presence of intermediaries. For instance, evidence by Ellison and Ellison (2009) suggests that a price-comparison website induces sellers to quote very low base prices and introduce high surcharges, undermining consumers' ability to do meaningful comparisons. More generally, firms may respond to price comparison aids with more obfuscation, again lowering the net effect on consumers.

Another, in our opinion more important, issue relates not to how firms react, but to the intermediaries themselves: whether and when profit-maximizing intermediaries guide consumers to correct choices rather than exploit consumers' fallacies just like producers do. Conceptually, a consumer—say, a retail investor—may search for two types of advice. First, she may want to decide which type of product—say, asset class—fits her

personal needs. Second, once she identifies a class of products, she may be looking to figure out which product is the best—say, which mutual fund among those investing in the US stock market to buy. Regarding the first problem, if fees are similar across products types an advisor has little reason to misguide the client—although also little reason to try hard to find the suitable product.

For the second class of problems, however, there is reason to believe that intermediaries are useless to harmful. These problems are especially relevant in retail finance, where an inferior product often amounts to a higher-fee version that provides the same service. Think of an investment advisor deciding between recommending an otherwise identical high- or low-fee fund, or a broker deciding between recommending one of two mortgages—a cheap one or a non-standard one that in expectation is more costly to the consumer. Murooka (2015) shows that competing intermediaries fail to educate consumers about very deceptive products, and their presence actually increases prices (see Section 4.1). This prediction is roughly consistent with the findings of an audit study by Mullainathan et al. (2011): if anything, advice exacerbates consumers' biases by encouraging the chasing of returns and investing in actively managed funds.

Research on other types of advice largely supports the above pessimistic conclusions. In Armstrong and Zhou's (2011) commission model, upstream firms sell a homogenous product to a population of naive and fully informed (sophisticated) consumers. Each upstream firm sets prices and a commission for promoting its product to consumers. Naive consumers visit an intermediary and credulously follow this intermediary's recommendation. Being unable to affect the behavior of informed consumers, an intermediary always promotes the product with the highest commission. The higher-price firm, therefore, only makes a sale if it offers a higher commission and the intermediary steers consumers to it. This means that a firm paying a low commission earns money only from informed consumers and hence prices aggressively to attract them, while a firm offering a very high commission is confident that the intermediary will recommend it to uninformed consumers, so it sets high prices. In the unique mixed-strategy equilibrium, therefore, prices are positively correlated with commissions, and naive consumers are recommended the higher-priced product.

The above problems would not arise if intermediaries were paid directly by consumers, and not through commissions for making sales. But Inderst and Ottaviani (2012b) show that when consumers ignore that commissions influence advice, then the fee-based business model does not emerge in the market.[58] If the advisor received a fixed fee, then a seller could raise the price of its product and the advisor reduce her fee, keeping the money that they receive jointly unchanged and hence a naive consumer indifferent. The producer can then use the increased revenue to pay a commission to

[58] See Inderst and Ottaviani (2012a) for a discussion of different policy interventions in the market for financial advice that also covers the case of naive consumers taking advice at face value.

the advisor, which induces the advisor to steer more naive consumers to it. Consumer naivete, hence, can explain the prevalence of a commission-based advice model.

7.2 Soft paternalism

Given that we cannot count on markets to supply the right information or products to naive consumers, researchers have proposed market interventions to improve outcomes. A large part of the literature emphasizes soft-paternalistic interventions. While different specific principles for soft paternalism have been proposed, the common thread is clear: we should help consumers without interfering too harshly in their environment. A policy satisfies the principle of libertarian paternalism (Thaler and Sunstein, 2003; Thaler and Benartzi, 2004; Thaler and Sunstein, 2008) if it induces better decisions by consumers who make mistakes, but it does not much restrict their freedom to choose as they would without the policy. And a policy satisfies asymmetric paternalism (Camerer et al., 2003) if it helps consumers who make mistakes, but it does not much harm consumers who are choosing optimally.

While non-controversial when used correctly, the application of these principles has ignored basic industrial-organization insights to an extent that we find problematic. When asking whether a policy is libertarian paternalistic or asymmetrically paternalistic, researchers and policymakers often assume that the supply in the marketplace remains unchanged, and ask whether consumers will be able to choose better. But industrial organization (indeed, equilibrium thinking) tells us that if we change the behavior of a significant share of consumers—the very aim of the policy—then supply is likely to change as well. When that happens, there is little hope for the policy to remain libertarian paternalistic, and all too often the policy cannot be asymmetrically paternalistic either. As a simple example, in any market model in which there is a cross-subsidy similar to the one analyzed in Section 2.3, a reduction in consumer mistakes makes sophisticated consumers worse off.

The same point applies to a common libertarian-paternalistic tool, defaults. The appeal is obvious: if a policymaker has a good idea about which product is likely to benefit consumers, it can nudge them toward choosing that product through defaults, with the confident knowledge that defaults have a major effect on behavior. But to convincingly evaluate the welfare implications, we need to model how libertarian-paternalistic interventions affect consumers' decisionmaking and, ideally, the firms' ability to educate or confuse as well as their contract offers (Spiegler, 2015). While we are still limited here by a lack of understanding of what exact psychological mechanisms drive default effects in different environments and, more generally, what consumers understand, Spiegler (2015) highlights in a set of examples that choice-architecture-style interventions can plausibly backfire when taking firms' responses into account.

Given the above recognition, a natural question arises: in what types of environments is it appropriate to evaluate soft-paternalistic interventions by assuming that the

market environment is fixed? One type of environment could be markets in which the supply is determined by the policymaker, for instance through a state monopoly. But even in this case, the change in consumer behavior induced by an intervention typically has budgetary implications. Another type of environment could be markets in which the other side is largely inelastic, and for some reason prices do not respond to an intervention. Here, interventions affecting organ donations is the best example: the demand for organs is largely independent of supply, and there is no price mechanism in place to clear the market. But these situations are exactly the types of environments in which industrial-organization analysis is unnecessary, so we do not discuss them further in this review. A further example may be situations in which naive and sophisticated consumers self-select, and sophisticated consumers are served by a competitive supply that is perfectly elastic—as in the mutual-fund industry model of Heidhues et al. (2017) in which sophisticated consumers buy competitively-supplied low-fee index funds and naive consumers buy high-fee managed funds. An intervention that induces naive consumers to take the high fees into account, or select the competitive index funds for other reasons, may qualify as soft-paternalistic.[59]

A second, interrelated, question is why a policymaker would want to restrict herself to libertarian or asymmetric paternalism and refrain from other interventions. In most classical industrial-organization papers, policies are analyzed from a total-welfare or perhaps consumer-welfare perspective, implicitly relying on potential Pareto improvements. It is unclear to us why we should shy away from a regulatory intervention in the banking market that would help naive consumers avoid overdraft fees, simply because it reduces the cross-subsidy to more sophisticated consumers (especially in reverse-Robin-Hood-like equilibria in which the poor cross-subsidize the rich). At least conceptually, a more promising approach to us would specify a welfare function to be maximized, and then look at the optimal regulation that achieves such a goal. In line with this idea, from now on we evaluate interventions from the classical vantage points of overall efficiency and distribution, not from the perspective of whether they satisfy principles of soft paternalism.

7.3 Disclosure and consumer education

The most obvious type of intervention aimed at mitigating the welfare-decreasing effects of consumer naivete targets the naivete itself. In particular, many researchers and policymakers have proposed that we might be able to induce more sophisticated behavior for at least a part of the population through education or improved disclosure. In an attempt to help boundedly rational consumers to choose between different complicated tariffs by different suppliers, for instance, Ofgem proposed requiring energy suppliers

[59] In practice, if the intervention induces naive consumers to select index funds without comprehending the reason why, of course, it may give rise to high-fee or otherwise exploitative index funds.

to display a "tariff comparison rate" similar to the APR for loans (Office of Gas and Electricity Markets, 2013). From many perspectives, education is uncontroversial: it can help not only naive consumers, but also consumers who are just uninformed in the classical sense; it is obviously soft paternalistic; and it is often easily accepted by firms. Yet a number of limitations to education have been identified in the literature.

Effectiveness of education. An immediate concern is whether the education is effective in changing consumers' understanding of the product or market in question. As we have emphasized, the fact that a product feature is disclosed does not in itself mean that consumers become aware of it and take it into account—this is ultimately an empirical question. Indeed, much of the evidence reviewed in the chapter Behavioral Household Finance of this handbook suggests that education campaigns often have little to no effect.[60]

Unfortunately, there is very little academic research on what kinds of education might be helpful to consumers, and we view this as the greatest gap in the literature on the topic. A notable exception is Bar-Gill and Ferrari (2010). They point out that existing disclosure requirements in the US and the EU almost exclusively focus on attributes of the product or contract offered. This kind of disclosure is not helpful for consumers who mispredict their own product use, a frequent mistake in the applications discussed in Sections 2 through 4. For this reason, Bar-Gill and Ferrari (2010) propose extending disclosure requirements to "product-use" information, such as the average monthly payment consumers make for mortgages or the average amount consumers pay in late fees for credit cards. Even better, they argue, is to require firms to disclose individual-level usage information to consumers in markets in which this information is collected anyhow. For example, a credit-card company may be required to disclose how often the individual consumer has paid late, which hopefully helps overcoming "better-than-average" effects.

Kiss (2014) documents that advertising campaigns can be effective in inducing consumers to overcome inertia—a common concern among regulators of gas, electricity, banking, health insurance markets—and increase switching rates and thereby competition. Kiss exploits a regulatory change in the Hungarian market for auto liability insurance. Prior to 2010, all insurance contracts ended with the calendar year and consumers could only switch in the month of November. This induced widespread advertising campaigns as well as newspaper articles comparing insurance options. After the regulatory change, all insurance contracts last a year, implying that they end yearly with the car purchase anniversary. Kiss finds a significant difference in switching rates between new-regime drivers who bought their car in the middle of the year and those

[60] Furthermore, even if disclosure is effective in a particular instance, it requires consumer attention, which is a limited resource. We return to this issue in the next subsection.

who bought it around January 1. As consumers are otherwise similar, Kiss argues that the increased switching rate of the latter consumers follows from increased attention to the switching decision following the advertising and information campaign.

Negative side-effects of education. Beyond the question of whether and when education is effective in reducing naivete, a number of authors have highlighted ways in which (effective) education/disclosure can have unintended negative consequences when naive consumers are present. While most of the mechanisms—especially those centered on the naive consumers "left behind" by education that is only partially effective—are compelling and potentially important, we think of them as making the case for additional interventions rather than as making the case that we should not attempt education.

A simple example of a negative side-effect of education is immediate from Eq. (2) and the discussion following it, and is analyzed in more detail by Kosfeld and Schüwer (2017). Suppose that we are in a market with a sophisticated-side distortion, and a regulator concerned with consumer naivete implements a successful education campaign that decreases the share α of naive consumers. Holding firms' offers fixed, this is clearly Pareto-improving (and hence asymmetrically paternalistic) and libertarian: educated naive consumers are strictly better off, while formerly sophisticated consumers and uneducated naive consumers are equally well off. But it is also clear that in equilibrium the intervention cannot always be Pareto-improving simply because it may lower total welfare. With more consumers being sophisticated, more consumers engage in inefficient avoidance behavior, so if the effect on the additional price is sufficiently small, lower total welfare results.[61] The intuition is easiest to see in a competitive market: given that fewer consumers now pay the additional price, firms must raise prices to break even, making formerly sophisticated consumers as well as uneducated naive consumers worse off.[62]

A host of more subtle reactions by firms can also have negative effects. Suppose disclosure regulation facilitates comparing products for sophisticated consumers and hence makes this market segment more competitive. Firms respond by refocusing their business model on earning money from naive consumers' mistakes (Murooka and Schwarz, 2018), potentially increasing exploitation distortions and decreasing consumer utility.

Similarly, education by a regulator can lower education—or increase obfuscation—by firms (e.g., Carlin and Manso, 2011; Piccione and Spiegler, 2012). To illustrate this logic in a particular situation, consider the bare-bones model used in Section 4 with an education-favoring technology and a binding price floor. Consider a disclosure—or any

[61] This logic makes clear that with naive-side or homogenous distortions the same concern does not arise.

[62] As a less stark instance of the same issue, Kamenica et al. (2011) provide a simple formal example in which educating consumers benefits them taking prices as given, but once the response of firms is taken into account, leaves them exactly as well off as before.

other—regulation that is effective in reducing the maximal unexpected payments (a_{max}) a firm can collect from naive consumers. Suppose that consumers are heterogenous in their valuations, with a fraction $\kappa \in (0, 1)$ of consumers having $v \geq \underline{f} + a_{max}$, and another fraction having a sufficiently low v for them to buy in equilibrium only if an obscure market obtains ($v \lesssim \underline{f}$). Then, the higher is a_{max}, the greater is a firm's incentive to educate consumers, and hence the more transparent the market tends to be. Intuitively, the proportional increase in profits from optimally undercutting the competitor's total price is increasing in that price.[63] Attempts to make the market more transparent by lowering a_{max} can therefore backfire by leading firms to obscure the market.

Another reason to be skeptical of the universal desirability of educating or debiasing consumers is based on the simple observation that consumer mistakes can mitigate other distortions. Consider an insurance market. If all consumers are rational, then adverse selection can be severe and can lead to highly inefficient outcomes. If some consumers are naive—specifically, if some healthy consumers choose expensive full coverage by mistake—then adverse selection is less severe, increasing welfare (Handel, 2013; Handel et al., 2016). This is not a mere theoretical possibility: Handel (2013) estimates that if consumers chose optimally, the welfare loss from adverse selection in the US health insurance market would double.[64]

A related observation applies in the attention model of Grubb (2015a), which we introduced in Section 3.1. Recall that Grubb analyzes services (such as mobile-phone calls or bank-account transactions) whose marginal price depends on consumers' past behavior, and because consumers may not recall their past behavior, they may not know their marginal price. One is then tempted to argue that consumers would be better off being notified when they approach or exceed their included allowance. Based on a dynamic structural model of the US mobile-phone industry between 2002 and 2004, Grubb and Osborne (2015) estimate that holding firms' offers fixed, consumers benefit substantially from such notifications. But once the equilibrium response from firms is taken into account, notifications lower welfare by $26 and consumer surplus by $33 per person per year. The reason is simple: because consumer inattention facilitates efficient screening by firms, disclosing marginal prices lowers welfare. Importantly, however, this point is

[63] Specifically, suppose the competitor's total price is $p_1 = \underline{f} + a_{max}$ in a candidate equilibrium without education. If the firm does not educate and charges the candidate equilibrium price, she earns $(p_1 - c)/2$. If the firm deviates and educates, she faces the demand $D_0(p_0, p_1) = \kappa(t + p_1 - p_0)/2t$ so given the constant marginal cost c her optimal deviation price is $p_0 = (t + p_1 + c)/2$, earning deviation profits of $(\kappa/8t)(t + p_1 - c)^2$. Differentiating the ratio between deviation and candidate equilibrium profits $[\kappa(t + p_1 - c)^2]/[4t(p_1 - c)]$ with respect to p_1 shows that is increasing in p_1 in the range where $p_1 > t + c$ (which holds since the price floor is binding in the candidate equilibrium). Hence, the incentive to deviate from a candidate shrouding equilibrium in which $p_1 = \underline{f} + a_{max}$ is increasing in a_{max}. Thus, for a given κ firms educate if a_{max} is large enough.

[64] See the chapter titled "Behavioral Economics and Health Care Markets" that is in preparation for Volume 2 of this handbook.

limited to rationally inattentive consumers—consumers who may not remember their past usage, but have rational expectations regarding their probability of running into a high fee. If consumers instead underestimate this probability, then (as also in Armstrong and Vickers, 2012) notifications can easily increase consumer and total welfare.

Problems regarding education can also arise from the type of education a regulator is able to provide. In an environment in which disclosing everything is ineffective, a regulator must simplify. But when consumers base their decisions on such simplified information, firms have an incentive to game the system. Duarte and Hastings (2012) empirically document this effect in the Mexican market for retirement financing. Due to heavy regulation, funds in this market choose very similar investment strategies, so consumers should select between funds based only on fees. But because firms could charge both a load and a management fee, investors—presumably unable to integrate the two fees into a single relevant price—initially reacted only little to the fees. As a result, funds charged high fees. To improve the market outcome for consumers, the regulator created an index that combined the two fees, with the explicit aim of facilitating comparison between funds. The index being salient, workers largely based their decisions on it. The firms readjusted their fee structure to optimally exploit the index, managing to increase revenues while lowering their index at the same time. The funds' response lowered the gain from the information-provision policy, and lead to a disadvantageous redistribution from low- to high-income consumers.

It is worth comparing the Mexican intervention with the intervention against drip pricing by airlines in the EU. Article 23 of Regulation (EC) No 1008/2008 requires that the displayed price must include "all applicable taxes, and charges, surcharges and fees which are unavoidable." While both interventions aim to simplify product comparison by introducing one relevant price measure in a market with multidimensional prices, economic logic suggests the intervention will work better for airlines. In the case of airlines, the different fees in question apply equally to all consumers, so integrating them into the headline price creates a price measure that accurately reflects what everyone will pay. Hence, if consumers are sensitive to the headline price, price competition results and benefits all consumers. In the case of funds, however, the load and the management fee do not apply equally to all consumers—the management fee is more important for investors with a long investment horizon—so a single index cannot capture the price for all consumers, and introduces a scope for firms to manipulate the index. Even if the index increases price competition, therefore, it does not necessarily benefit all consumers.

7.4 Regulating contracts or firm conduct

In addition to soft interventions, another approach to combat the effects of consumer mistakes is regulating the products or contracts themselves. The aim of these interventions is eliminating product or contract features that are being used to exploit naive

consumers, and that are unlikely to serve useful economic purposes. Many researchers and policymakers seem to prefer considering hard interventions only after potential soft interventions have been exhausted. We are unaware of any logical or empirical argument for taking this approach, and believe that regulations should be considered in parallel with soft interventions.

To illustrate the potential of regulation, consider a type of widespread hard intervention that is not typically discussed in economics: safety regulations. As Bar-Gill and Warren (2008) point out, extensive safety regulations are ubiquitous for products ranging from toasters to car seats. Yet in a model with rational consumers who understand all disclosed information, the case for safety regulations appears weak: rather than restricting what products can be sold, the regulator can simply require disclosure of risks, so that consumers can make their own decisions regarding what to buy. Yet safety regulations make perfect sense when viewed from the perspective of trying to reduce the scope for consumer mistakes in purchases. Most consumers who would purchase a crib that puts babies' lives in danger would do so by mistake rather than after a calculated tradeoff between price and safety. Banning dangerous cribs eliminates the possibility of mistakes and creates almost no distortion from consumers who would rationally prefer less safe cribs. Of course, defining a safe crib is highly product-specific, and firms looking to skimp on costs will try to circumvent the definition. Even so, safety regulations are effective enough that consumers can shop without having to worry that a crib will collapse.

While similarly far-reaching regulations do not exist for many types of contracts, the same case as for safety regulations can be made: if a contract feature is likely to induce many mistakes and has little economic purpose, then banning it is welfare-increasing. As in the case of physical products, such regulation is difficult, likely to be market-specific, and firms will have strong incentives to circumvent it. These considerations must be taken into account when designing the regulation, but they are not reasons to foreclose considering regulations altogether. In the rest of the subsection, we discuss the potential and pitfalls of some regulations in specific contexts.

Regulating exploitative features. When recognizing the presence of naive consumers, an obvious potential regulation is lowering the contingent charges that naive consumers may ignore. For instance, the OFT took enforcement action against the pricing practices of gyms in the UK, which resulted in some terms, such as minimum contract lengths and high termination fees, being deemed unfair and contrary to good faith (Nuñez, 2017). A more complex example is the regulation of debt settlement services in the US.[65] In a common debt settlement program prior to the regulation, consumers were required to make payments to be used both for the provider's fees and

[65] See https://www.federalregister.gov/documents/2010/08/10/2010-19412/telemarketing-sales-rule.

for eventually settling a portion of the consumer's debt. Once sufficient funds were accumulated, the provider began negotiations with the creditor. Consumers who did not complete the—often multi-year—program forfeited the provider's fees without receiving any service. Charging for services that have not yet been rendered is therefore no longer permitted.

These types of interventions, equivalent to lowering a_{max} in our bare-bones model, are discussed in the context of the UK market for current accounts by Armstrong and Vickers (2012). Because the additional price facilitates a cross-subsidy from naive to sophisticated consumers, lowering it through regulation benefits naive consumers and harms sophisticated consumers. This point must be qualified when (as in Heidhues et al., 2017) there is a binding floor on the anticipated price. In the range where the price floor remains binding, a decrease in a_{max} does not affect the anticipated price, so it benefits naive consumers without affecting sophisticated consumers. Consistent with this prediction, Bar-Gill and Bubb (2012) and Agarwal et al. (2015) find evidence suggesting 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.[66] In addition, in a model where sophisticated consumers can exert socially inefficient effort to avoid the additional price, a decrease in a_{max} to a level where sophisticated consumers no longer exert the effort increases overall welfare.[67] But Heidhues et al. (2016) point out a potential problem with this kind of regulation: it often increases firms' incentive to invent new hidden fees, lowering the net effect of the policy.

In the credit-market model of Heidhues and Kőszegi (2010) discussed in Section 2.5, naive borrowers underestimate how much they will pay in interest and fees and therefore underestimate the cost of credit, leading them to overborrow. To protect borrowers, regulations restrict practices generating large penalties: in July 2008 the Federal Reserve Board severely limited the use of prepayment penalties, and the Credit CARD Act of 2009 prohibits the use of interest charges for partial balances the consumer has paid off, and restricts fees in other ways. Our model predicts that because these and other regulations limiting unexpected payments can lower consumers' mispredictions, they can increase welfare.

Multilevel marketing schemes that use independent sales representatives who are paid for both selling a product as well as for acquiring new independent sales representatives are widespread (according to Antler, 2018, the size of the US multilevel marketing industry exceeds $35 billion). They are especially controversial if agents are not just rewarded for the sales agents they recruit themselves but also for the recruits attracted

[66] More generally, the pass-through of the revenues of the additional price to the anticipated one depends on demand and supply-elasticities (for a discussion thereof see Grubb, 2015c).

[67] For further potential welfare effects of regulating the additional price, see our discussion of the participation and exploitation distortions in Section 2.4.

by their recruits, etc.[68] Indeed, a system of (high) rewards for generating a downline resembles pyramid scams, which are illegal in most countries. Antler (2018) develops a behavioral contracting model to distinguish between exploitative pyramid schemes and incentive systems for rational agents, which suggests that consumer protection agencies may indeed want to rule out rewards to a sales agent for the downline she generates.[69]

Heidhues et al. (2018) highlight the pro-competitive implication of regulating secondary contract features—such as safety aspects of a product, contract clauses, or additional price components—when consumers' attention is limited. In their basic model, each firm's contract offer consists of a headline price as well as an additional price. Consumers initially see the headline price of a randomly chosen firm, and then decide whether to spend their available attention on studying this firm's contract offer—in which case they get to know the additional price and when it applies—or browsing another firm—in which case they learn about that firm's existence and its headline price. Contract regulations that limit additional prices or set default conditions under which they do not apply intensify competition and thereby increase consumer welfare in this benchmark and many related environments. The underlying reason is straightforward: freeing consumers from having to study the regulated features allows them to spend attention on comparing products instead, increasing competition. When allowing for multiple markets on which the consumer can expend a given amount of attention, they show that the benefits to consumers from regulating a given market may occur in other markets, and that these benefits are highly non-linear: once the regulation covers sufficiently many markets, a strong pro-competitive effect kicks in.

But as Heidhues et al. (2018) emphasize, one must be careful about what price-related aspect of a product one regulates. Building on and extending Fershtman and Fishman (1994), Armstrong et al. (2009) consider a search model in which consumers can exert effort to become informed about the best deals in the market. A price cap restricts equilibrium price dispersion and thereby lowers consumers' incentive to become informed. As a result, such a policy can reduce price competition and increase the average price consumers pay.[70] A regulator must therefore restrict only the secondary features of products that take advantage of consumer naivete or limited attention, and

[68] See, for example, the warnings by the Federal Trade Commission: https://www.ftc.gov/tips-advice/business-center/guidance/multilevel-marketing (accessed on May 5, 2018).

[69] Antler establishes that if a firm sells a valuable good through an optimal incentive scheme with rational agents, then it does not rely on rewarding agents for the downline they generate. The same, however, is no longer true with plausible forms of agent naivete. In particular, Antler solves for an analogy-based expectation equilibrium (Jehiel, 2005) in which agents ignore that it becomes more difficult to sign up further agents later in the game. Similarly to other contracting models with naive agents, it is optimal for a firm to reward outcomes whose likelihood the agent overestimates.

[70] Similarly, the authors demonstrate that a "do-not-call" list can have a detrimental effect if it reduces consumers' price knowledge. While it can be privately optimal to join such a list, uninformed consumers reduce price competition between firms and so joining the list generates a negative externality.

not interfere with the core price mechanism in the market.[71] In practice, it might often be difficult to precisely distinguish the two.

Changing sellers' incentives. Some interventions proposed in the literature work by changing the incentives of firms or agents who have contact with consumers at the point of sale. The goal of such interventions is to reduce the benefit the seller derives from exploiting consumer mistakes. In the context of financial advice, the model by Murooka (2015) discussed in Section 4.1 implies that making commissions independent of the type of product leads intermediaries to sell better products to consumers. With uniform commissions, deceptive firms cannot use the high profits they make from exploiting naive consumers to set high commissions, and therefore cannot squeeze transparent products out of the market. Indeed, recent US regulation allows a broker's fee for a mortgage to depend only on the size and not the terms or conditions of the mortgage.[72]

A completely different approach is to lower the profits a firm can make from the additional price. As a case in point, suppose the additional price is generated from selling an add-on good or service (see Section 2 for examples). Then, if it is feasible to induce competition in the add-on market, this will eliminate the profits a firm can earn from naive consumers, and increase the anticipated price, reducing the inefficiency. As discussed and analyzed in Michel (2017), this idea underlies recent regulations of extended warranties in the UK by the Office of Fair Trading that aim to increase competition for these warranties at the point of sale.[73]

Korobkin (2003) proposes that many exploitative contract features can be policed by ex-post judicial review under the unconscionability doctrine, a legal doctrine that invalidates contract terms over which a party had no effective choice. Korobkin argues that it is logically impossible for a consumer to have effective choice over a contract feature she does not understand, so that this situation should be covered under unconscionability. Nevertheless, this is not how courts have interpreted the unconscionability doctrine so far. Similarly, Bar-Gill (2004) argues that the penalty doctrine, which precludes damages beyond the harm or reasonable ex-ante estimate of the harm caused to the other party, is applicable to some credit-card fees. For instance, a $30 late fee for paying a $10 balance a few days late is clearly illegal under the doctrine. Once again, however, courts have been reluctant to invoke the penalty doctrine for many types of contract terms. While

[71] Heidhues et al. (2018) argue that the European Union's principle on unfair contract terms—despite typically being motivated solely on the basis of fairness concerns—broadly matches the model's policy recommendation. Similarly, some existing safety regulations that do not regulate the price or functionality of the core product are in line with the model's policy prescription.

[72] See 12 CFR §226.36.

[73] In contrast, Michel (2017) argues that a minimum quality standard is less effective than inducing aftermarket competition, and may even lower consumer surplus.

these proposals are promising, it would be useful to evaluate the potential of ex-post judicial review in the context of economic theories in which consumers make mistakes.

Political economy of regulation. Beyond the difficulty of finding welfare-increasing regulations and the potential negative side-effects, Warren and Wood (2014) highlight a general political-economy problem with many hard regulations: even if they help naive consumers, they tend to be disliked by all—i.e., both naive and sophisticated—consumers. Although Warren and Wood's result is much more general, the main idea is easily illustrated in the bare-bones model analyzed in Section 2.3, in which sophisticated consumers benefit from receiving a cross-subsidy from naive consumers. These sophisticated consumers, then, are hurt if some contract regulation reduces the exploitation of naive consumers and should vote against it. Naive consumers, who erroneously believe themselves to be sophisticated, will thus also not be in favor of such a regulation.

Nevertheless, several forces outside Warren and Wood's model make it possible to enact regulations that reduce consumer mistakes. First, regulators may not ask for voter approval on many interventions they consider. Second, consumers who are not fully sophisticated may underappreciate the equilibrium effects of regulation—just as many economists may before becoming familiar with the relevant research—so they might not understand that the regulations in question will hurt sophisticated consumers. Third, citizens concerned about general welfare may appreciate regulations that do not necessarily benefit them.

7.5 Modifying classical policy approaches and recommendations

Questioning inferences about anti-competitive practices. Ronald Coase made a cogent observation that is relevant both for the understanding of markets and for the understanding of industrial organization as a science: "[o]ne important result of [the] preoccupation with the monopoly problem is that if an economist finds something—a business practice of one sort or other—that he does not understand, he looks for a monopoly explanation. And as we are very ignorant . . . the number of ununderstandable practices tends to be very large" (Coase, 1988, p. 67). Research in behavioral industrial organization lends support to a variant of Coase's judgment. Namely, the unrealistically narrow view of consumers and managers in the classical industrial-organization paradigm can lead to a misinterpretation of market phenomena, and thereby to a misguided call for market intervention. Most importantly, a number of puzzling patterns in firm behavior have been attributed to harmful practices such as collusion or predation, when in fact plausible explanations based on richer models of consumer behavior are possible. Although we are unaware of careful evidence to distinguish the mechanisms at play—and more research on how to do so would be extremely valuable—we provide two potentially important examples.

First, models of consumer loss aversion (Heidhues and Kőszegi, 2008) provide an explanation for the lack of price variation that is different from theories of collusive behavior typically invoked to explain the same pricing practices (see Athey et al., 2004, and the more informal arguments preceding it).

Our second example is motivated by the model of unplanned purchases by Johnson (2017). Suppose consumers with horizontally differentiated preferences over stores purchase two products—milk and soap—with probability one, but ex ante they erroneously believe that they will purchase soap only with probability 1/2. Denoting a consumer's value from the two products by $v_m > 0$ and $v_s > 0$ and prices by p_m and p_s, a consumer's anticipated utility of visiting a store, gross of transportation costs, is $v_m - p_m + (1/2)(v_s - p_s)$. Hence, firms always set $p_s = v_s$; otherwise, a firm could raise p_s and lower p_m by the same amount, keeping revenue constant but increasing consumers' perceived utility, and thus demand. Hence, goods consumers anticipate to buy for certain—staple goods—have lower mark-ups and in sufficiently competitive settings are priced below cost. Now suppose that one firm—say, a corner store—can only stock one of the two products. Then, it will tend to stock the staple product, because the higher is consumers' anticipated purchase probability, the higher is consumers' anticipated surplus. Therefore, small stores stock exactly those goods for which large stores have low (possibly negative) markups, providing a new explanation for a practice that is typically interpreted as predatory pricing.

Furthermore, if firms make mistakes or follow goals other than profit maximization, we need to be careful when inferring unobservable variables from firms' behavior. For instance, consider again predatory pricing—reducing one's current price to increase future market power or even drive a rival out of the market. The classic profit-maximization hypothesis suggests that the possibility of predation can be ruled out unless it is feasible for the firm to recoup the current costs of aggressive pricing by the future benefits of market power, and legal practice is based on this (Bolton et al., 2000). But in as much as firms' managers are driven by vengeance or relative profits, such conclusions need not hold. Similarly, if small firms make less strategic supply decisions—as Hortaçsu and Puller (2008) estimated for the Texas electricity market—this can lead to significant welfare losses, suggesting that mergers may increase efficiency by increasing firms' strategic sophistication. Future research on behavioral firms will hopefully help in predicting when such concerns are important, and how they could be addressed.

Insurance. In addition to questioning the premise behind some classical approaches, the literature in behavioral industrial organization offers a few specific modifications to classical policy insights. An important example is insurance markets. It is well-known that in Rothschild and Stiglitz's (1976) model of the insurance market, low-risk agents receive cheap insurance, but this insurance is partial; otherwise, high-risk types would prefer to purchase the same insurance, and then the price could not be kept low. Because

some consumers are underinsured, a government policy of mandatory insurance can be Pareto-improving. Sandroni and Squintani (2007) show that the case can be weaker with overconfident consumers (consumers who are high-risk but believe themselves to be low-risk). As explained in Section 3.1, with overconfident agents the group of low-risk and overconfident are offered insurance they believe is actuarially unfair, so that they may prefer no insurance, or they may find the amount of insurance they can purchase in the market sufficient. In this case, a mandatory expansion in insurance coverage is not Pareto-improving. A similar effect occurs in Schumacher (2016).

Relatedly, Handel's (2013) work mentioned above demonstrates that consumer inertia renders adverse selection less important than in a model with rational consumers. The obvious implication is that steps a social planner can take to reduce adverse selection may not be as important as we would expect with rational consumers. Looking at this issue more carefully, Handel et al. (2016) show that the effect of risk adjustment (a policy that combats adverse selection by making its effect on insurer profits less severe) depends on the quality of consumer choices, and the policies in place to improve consumer choices. In particular, since improving consumer choices can exacerbate adverse selection, it is more likely to improve welfare if risk adjustment is also in place. Hence, the two types of policies are complementary.

Privacy. Although the research is in its infancy, the existing literature suggests that the case for protecting private information is often, but now always, stronger when consumers make mistakes than when they do not. Heidhues and Kőszegi (2017) demonstrate that with homogenous distortions, the implications of firms knowing more about consumers are often opposite with naivete-based discrimination than with classical preference-based discrimination (Section 3.2). Presumably, then, the implications for protecting privacy are opposite as well. But the same is not the case for sophisticated-side and naive-side distortions. And Hoffmann et al. (2014, described in Section 3.3) find that naivete about targeted advertisement can lower welfare, but only in less competitive environments where firms can price discriminate.

Benefits from product variety. In classic industrial organization, to infer the benefits from variety, researchers estimate a demand system and calculate the resulting consumer surplus. If consumer choices are partly based on mistakes or consumers only consider a subset of all available goods, however, then the inference is misleading and researchers are likely to overestimate the benefit from product variety (and product innovation). This is easiest to see in the extreme case in which all goods are homogenous but consumers have problems comparing prices or products, so that they behave as if products were differentiated.

REFERENCES

Agarwal, Sumit, Chomsisengphet, Souphala, Mahoney, Neale, Stroebel, Johannes, 2015. Regulating consumer financial products: evidence from credit cards. Quarterly Journal of Economics 130 (1), 111–164. Previously NBER Working Paper #19484.

Akerlof, George, 1970. The market for 'lemons': quality uncertainty and the market mechanism. Quarterly Journal of Economics 84 (3), 488–500.

Amador, Manuel, Werning, Iván, Angeletos, George-Marios, 2006. Commitment vs. flexibility. Econometrica 74 (2), 365–396.

Amaldoss, Wilfred, Jain, Sanjay, 2005. Pricing of conspicuous goods: a competitive analysis of social effects. Journal of Marketing Research 42 (1), 30–42.

Amaldoss, Wilfred, Jain, Sanjay, 2008. Trading up: a strategic analysis of reference group effects. Marketing Science 27 (5), 932–942.

Antler, Yair, 2018. Multilevel Marketing: Pyramid-Shaped Schemes or Exploitative Scams. Working Paper.

Apffelstaedt, Arno, Mechtenberg, Lydia, 2017. Competition over Context-Sensitive Consumers. Working Paper.

Armstrong, Mark, 2015. Search and ripoff externalities. Review of Industrial Organization 47 (3), 273–302.

Armstrong, Mark, Huck, Steffen, 2010. Behavioral economics as applied to firms: a primer. Competition Policy International 6 (1), 3–45.

Armstrong, Mark, Vickers, John, 2012. Consumer protection and contingent charges. Journal of Economic Literature 50 (2), 477–493.

Armstrong, Mark, Zhou, Jidong, 2011. Paying for prominence. Economic Journal 121 (556), F368–F395.

Armstrong, Mark, Vickers, John, Zhou, Jidong, 2009. Consumer protection and the incentive to become informed. Journal of the European Economic Association 7 (2–3), 399–410.

Artinger, Sabrina, Powell, Thomas C., 2016. Entrepreneurial failure: statistical and psychological explanations. Strategic Management Journal 37 (6), 1047–1064.

Astebro, Thomas, Herz, Holger, Nanda, Ramana, Weber, Roberto A., 2014. Seeking the roots of entrepreneurship: insights from behavioral economics. Journal of Economic Perspectives 28 (3), 49–70.

Athey, Susan, Bagwell, Kyle, Sanchirico, Chris, 2004. Collusion and price rigidity. Review of Economic Studies 71 (2), 317–349.

Ausubel, Lawrence M., 1991. The failure of competition in the credit card market. American Economic Review 81 (1), 50–81.

Axinn, Stephen M., 1983. A lawyer's response. Antitrust Law Journal 52 (3), 643–649.

Bagwell, Kyle, 2007. The Economic Analysis of Advertising. Handbook of Industrial Organization, vol. 3. Elsevier.

Bagwell, Laurie Simon, Bernheim, B. Douglas, 1996. Veblen effects in a theory of conspicuous consumption. American Economic Review 86 (3), 349–373.

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

Bar-Gill, Oren, 2009. The law, economics and psychology of subprime mortgage contracts. Cornell Law Review 94 (5).

Bar-Gill, Oren, Bubb, Ryan, 2012. Credit card pricing: the CARD act and beyond. Cornell Law Review 97.

Bar-Gill, Oren, Ferrari, Franco, 2010. Informing consumers about themselves. Erasmus Law Review 3 (2), 93–119.

Bar-Gill, Oren, Warren, Elizabeth, 2008. Making credit safer. University of Pennsylvania Law Review 157 (1), 1–101.

Baron, David P., 2009. A positive theory of moral management, social pressure, and corporate social performance. Journal of Economics & Management Strategy 18 (1), 7–43.

Bartoš, Vojtěch, Bauer, Michal, Chytilová, Julie, Matějka, Filip, 2016. Attention discrimination: theory and field experiments with monitoring information acquisition. American Economic Review 106 (6), 1437–1475.

Baumol, William J., 1958. On the theory of oligopoly. Economica 25 (99), 187–198.
Beck, Jonathan, 2004. Fixed, Focal, Fair? Book Prices with Optional Resale Price Maintenance. WZB Discussion Paper SP II 2004-15.
Becker, Gary S., Grossman, Michael, Murphy, Kevin M., 1994. An empirical analysis of cigarette addiction. American Economic Review 84 (3), 396–418.
Bénabou, Roland, Tirole, Jean, 2010. Individual and corporate social responsibility. Economica 77 (305), 1–19.
Benabou, Roland J., Tirole, Jean, 2016. Bonus culture: competitive pay, screening, and multitasking. Journal of Political Economy 124 (2), 305–370.
Bertrand, Marianne, Schoar, Antoinette, 2003. Managing with style: the effect of managers on firm policies. The Quarterly Journal of Economics 118 (4), 1169–1208.
Bloom, Nicholas, van Reenen, John, 2007. Measuring and explaining management practices across firms and countries. Quarterly Journal of Economics 122 (4), 1351–1408.
Bloom, Nicholas, Eifert, Benn, Mahajan, Aprajit, McKenzie, David, Roberts, John, 2013. Does management matter? Evidence from India. The Quarterly Journal of Economics 128 (1), 1–51.
Bolton, Patrick, Brodley, Joseph F., Riordan, Michael H., 2000. Predatory pricing: strategic theory and legal policy. Georgetown Law Journal 88 (8), 2239–2330.
Bordalo, Pedro, Gennaioli, Nicola, Shleifer, Andrei, 2013. Salience and consumer choice. Journal of Political Economy 121 (5), 803–843.
Bordalo, Pedro, Gennaioli, Nicola, Shleifer, Andrei, 2016. Competition for attention. Review of Economic Studies 83 (2), 481–513.
Borenstein, Severin, Mackie-Mason, Jeffrey K., Netz, Janet S., 2000. Exercising market power in proprietary aftermarkets. Journal of Economics & Management Strategy 9 (3), 157–188.
Braithwaite, Dorothea, 1928. The economic effects of advertisement. The Economic Journal 38 (149), 16–37.
Brown, Alexander L., Camerer, Colin F., Lovallo, Dan, 2012. To review or not to review? Limited strategic thinking at the movie box office. American Economic Journal: Microeconomics 4 (2), 1–26.
Bubb, Ryan, Kaufman, Alex, 2013. Consumer biases and mutual ownership. Journal of Public Economics 105 (C), 39–57.
Bundeskartellamt, 2017. Background Information on the Facebook Proceeding (19.12.2017). Technical Report 2017.
Bushong, Benjamin, Rabin, Matthew, Schwartzstein, Joshua, 2017. A Model of Relative Thinking. Working Paper.
Busse, Meghan R., Lacetera, Nicola, Pope, Devin G., Silva-Risso, Jorge, Sydnor, Justin R., 2013. Estimating the effect of salience in wholesale and retail car markets. American Economic Review 103 (3), 575–579.
Cabral, Luis, 2018. We're number 1: price wars for market share leadership. Management Science 64 (5), 1975–2471.
Camerer, Colin, Issacharoff, Samuel, Loewenstein, George, O'Donoghue, Ted, Rabin, Matthew, 2003. Regulation for conservatives: behavioral economics and the case for 'asymmetric paternalism'. University of Pennsylvania Law Review 151 (3), 1211–1254.
Camerer, Colin F., Ho, Tech H., Chong, Juin K., 2004. A cognitive hierarchy model of one-shot games. Quarterly Journal of Economics 119 (3), 861–898.
Capen, E.C., Clapp, R.V., Campbell, W.M., 1971. Competitive bidding in high-risk situations. Journal of Petroleum Technology 23, 641–653.
Carbajal, Juan Carlos, Ely, Jeffrey C., 2016. A model of price discrimination under loss aversion and state-contingent reference points. Theoretical Economics 11, 455–485.
Carbajal, Juan Carlos, Hall, Jonathan, Li, Hongyi, 2016. Inconspicuous Conspicuous Consumption. Working Paper.

Carlin, Bruce I., 2009. Strategic price complexity in retail financial markets. Journal of Financial Economics 91 (3), 278–287.
Carlin, Bruce Ian, Manso, Gustavo, 2011. Obfuscation, learning, and the evolution of investor sophistication. Review of Financial Studies 24 (3), 754–785.
Cherington, Paul T., Copeland, Morris A., 1925. The economics of advertising – discussion. American Economic Review 15 (1), 36–41.
Chetty, Raj, Looney, Adam, Kroft, Kory, 2009. Salience and taxation: theory and evidence. American Economic Review 99 (4), 1145–1177.
Chiappori, Pierre-Andre, Salanie, Bernard, 2000. Testing for asymmetric information in insurance markets. Journal of Political Economy 108 (1), 56–78.
Chioveanu, Ioana, Zhou, Jidong, 2013. Price competition with consumer confusion. Management Science 59 (11), 2450–2469.
Cho, Sungjin, Rust, John, 2010. The flat rental puzzle. Review of Economic Studies 77 (2), 560–594.
Coase, R.H., 1937. The nature of the firm. Economica 4 (16), 386–405.
Coase, Ronald Harry, 1988. The Firm, the Market, and the Law. The University of Chicago Press.
Competition Commission of the United Kingdom, 1994. The Supply of Recorded Music: A Report on the Supply in the UK of Prerecorded Compact Discs, Vinyl Discs and Tapes Containing Music. http://www.competition-commission.org.uk/rep_pub/reports/1994/356recordedmusic.htm#full.
Cooper, Arnold C., Woo, Carolyn Y., Dunkelberg, William C., 1998. Entrepreneurs' perceived chances for success. Journal of Business Venturing 3 (2), 97–108.
Covert, Thomas R., 2015. Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale. Working Paper.
Dahremöller, Carsten, 2013. Unshrouding for competitive advantage. Journal of Economics & Management Strategy 22 (3), 551–568.
De Clippel, Geoffroy, Eliaz, Kfir, Rozen, Kareen, 2014. Competing for consumer inattention. Journal of Political Economy 122 (6), 1203–1234.
de la Rosa, Leonidas Enrique, 2011. Overconfidence and moral hazard. Games and Economic Behavior 73 (2), 429–451.
de Meza, David, Reyniers, Diane, 2012. Every shroud has a silver lining: the visible benefits of hidden surcharges. Economics Letters 116 (2), 151–153.
de Roos, Nicolas, 2017. Collusion with limited product comparability.
DellaVigna, Stefano, Gentzkow, Matthew, 2017. Uniform Pricing in US Retail Chains. Working Paper.
DellaVigna, Stefano, Malmendier, Ulrike, 2004. Contract design and self-control: theory and evidence. Quarterly Journal of Economics 119 (2), 353–402.
Dertwinkel-Kalt, Markus, Köster, Mats, 2018. Salience and Online Sales: The Role of Brand Image Concerns. Working Paper.
Doraszelski, Ulrich, Lewis, Greg, Pakes, Ariel, 2018. Just starting out: learning and equilibrium in a new market. American Economic Review 108 (3), 565–615.
Duarte, Fabian, Hastings, Justine S., 2012. Fettered Consumers and Sophisticated Firms: Evidence from Mexico's Privatized Social Security Market. NBER Working Paper #18582.
Einav, Liran, Orbach, Barak Y., 2007. Uniform Prices for differentiated goods: the case of the movie-theater industry. International Review of Law and Economics 27 (2), 129–153.
Eliaz, Kfir, Spiegler, Ran, 2006. Contracting with diversely naive agents. Review of Economic Studies 73 (3), 689–714.
Eliaz, Kfir, Spiegler, Ran, 2008. Consumer optimism and price discrimination. Theoretical Economics 3 (4), 459–497.
Eliaz, Kfir, Spiegler, Ran, 2011a. Consideration sets and competitive marketing. Review of Economic Studies 78 (1), 235–262.
Eliaz, Kfir, Spiegler, Ran, 2011b. On the strategic use of attention grabbers. Theoretical Economics 6, 127–155.

Ellison, Glenn, 2005. A model of add-on pricing. Quarterly Journal of Economics 120 (2), 585–637.
Ellison, Glenn, 2006. Bounded rationality in industrial organization. In: Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, vol. 2, pp. 142–174.
Ellison, Glenn, Ellison, Sara Fisher, 2009. Search, obfuscation, and price elasticities on the internet. Econometrica 77 (2), 427–452.
Ellison, Glenn, Wolitzky, Alexander, 2012. A search cost model of obfuscation. RAND Journal of Economics 43 (3), 417–441.
Englmaier, Florian, Schmöller, Arno, Stowasser, Till, 2018. Price discontinuities in an online market for used cars. Management Science 64 (6), 2754–2766.
Esteban, Susanna, Miyagawa, Eiichi, 2006. Temptation, self-control, and competitive nonlinear pricing. Economics Letters 90 (3), 348–355.
Esteban, Susanna, Miyagawa, Eiichi, Shum, Matthew, 2007. Nonlinear pricing with self-control preferences. Journal of Economic Theory 135 (1), 306–338.
Farrell, Joseph, Klemperer, Paul, 2007. Coordination and lock-in: competition with switching costs and network effects. In: Armstrong, Mark, Porter, Rob (Eds.), Handbook of Industrial Organization, vol. 3. North-Holland.
Fershtman, Chaim, Fishman, Arthur, 1994. The 'perverse' effects of wage and price controls in search markets. European Economic Review 38 (5), 1099–1112.
Fershtman, Chaim, Judd, Kenneth L., 1987. Equilibrium incentives in oligopoly. American Economic Review 77 (5), 927–940.
Finkelstein, Amy, McGarry, Kathleen, 2006. Multiple dimensions of private information: evidence from the long-term care insurance market. American Economic Review 96 (4), 938–958.
Friedrichsen, Jana, 2017. Signals Sell: Optimal Product Lines when Consumers Differ both in Taste for Quality and Image. Working Paper.
Gabaix, Xavier, 2011. A Sparsity-Based Model of Bounded Rationality. Working paper. NYU Stern. http://pages.stern.nyu.edu/~xgabaix/papers/sparsebr.pdf.
Gabaix, Xavier, Laibson, David, 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. Quarterly Journal of Economics 121 (2), 505–540.
Galperti, Simone, 2015. Commitment, flexibility, and optimal screening of time inconsistency. Econometrica 83 (4), 1425–1465.
Gamp, Tobias, Krähmer, Daniel, 2017. Deceptive Products and Competition in Search Markets. Working Paper.
Glaeser, Edward L., Ujhelyi, Gergely, 2010. Regulating misinformation. Journal of Public Economics 94 (3), 247–257.
Goldfarb, Avi, Xiao, Mo, 2011. Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. American Economic Review 101 (7), 3130–3161.
Goldfarb, Avi, Xiao, Mo, 2017. Transitory Shocks, Limited Attention, and a Firm's Decision to Exit. Working Paper.
Goldfarb, Avi, Yang, Botao, 2009. Are all managers created equal? Journal of Marketing Research 46 (5), 612–622.
Gottlieb, Daniel, 2008. Competition over time-inconsistent consumers. Journal of Public Economic Theory 10 (4), 673–684.
Gottlieb, Daniel, Smetters, Kent, 2012. Narrow Framing and Life Insurance. NBER Working Paper #18601.
Grossman, Sanford J., 1981. The informational role of warranties and private disclosure about product quality. Journal of Law and Economics 24, 461–489.
Grossman, Sanford J., Hart, Oliver D., 1986. The costs and benefits of ownership: a theory of vertical and lateral integration. Journal of Political Economy 94 (4), 691–719.
Grubb, Michael D., 2009. Selling to overconfident consumers. American Economic Review 99 (5), 1770–1805.

Grubb, Michael D., 2015a. Consumer inattention and bill-shock regulation. Review of Economic Studies 82 (1), 219–257.
Grubb, Michael D., 2015b. Failing to choose the best price: theory, evidence, and policy. Review of Industrial Organization 47 (3), 303–340.
Grubb, Michael D., 2015c. Overconfident consumers in the marketplace. Journal of Economic Perspectives 29 (4), 9–36.
Grubb, Michael D., 2015d. Selling to Loss Averse Consumers: A Survey. Working Paper.
Grubb, Michael D., Osborne, Matthew, 2015. Cellular service demand: biased beliefs, learning, and bill shock. American Economic Review 105 (1), 234–271.
Gul, Faruk, Pesendorfer, Wolfgang, 2001. Temptation and self-control. Econometrica 69 (6), 1403–1435.
Gurun, Umit G., Matvos, Gregor, Seru, Amit, 2016. Advertising expensive mortgages. Journal of Finance 71 (5), 2371–2416.
Hahn, Jong-Hee, Kim, Jinwoo, Kim, Sang-Hyun, Lee, Jihong, 2018. Price discrimination with loss averse consumers. Economic Theory 65 (3), 681–728.
Hall, Robert E., 1997. The Inkjet Aftermarket: An Economic Analysis. Unpublished Manuscript. Stanford University.
Hall, Robert E., Woodward, Susan E., 2010. The burden of the nondiversifiable risk of entrepreneurship. American Economic Review 100 (3), 1163–1194.
Hamilton, Barton H., 2000. Does entrepreneurship pay? An empirical analysis of the returns to self-employment. Journal of Political Economy 108 (3), 604–631.
Handel, Benjamin R., 2013. Adverse selection and inertia in health insurance markets: when nudging hurts. American Economic Review 103 (7), 2643–2682.
Handel, Benjamin R., Kolstad, Jonathan T., Spinnewijn, Johannes, 2016. Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets. Working Paper.
Hanna, Rema, Mullainathan, Sendhil, Schwartzstein, Joshua, 2014. Learning through noticing: theory and evidence from a field experiment. Quarterly Journal of Economics 129 (3), 1311–1353.
Hart, Oliver, 2009. Hold-up, asset ownership, and reference points. Quarterly Journal of Economics 124 (1), 267–300.
Hart, Oliver, Holmstrom, Bengt, 2010. A theory of firm scope. Quarterly Journal of Economics 125 (2), 483–513.
Hart, Oliver, Moore, John, 1990. Property rights and the nature of the firm. Journal of Political Economy 98 (6), 1119–1158.
Hart, Oliver, Moore, John, 2008. Contracts as reference points. Quarterly Journal of Economics 123 (1), 1–48.
Hastings, Justine S., Hortaçsu, Ali, Syverson, Chad, 2017. Sales force and competition in financial product markets: the case of Mexico's social security privatization. Econometrica 85 (6), 1723–1761.
Hefti, Andreas, 2016. Limited Attention, Competition and Welfare. Working Paper.
Hefti, Andreas, Liu, Shuo, 2016. Targeted Information and Limited Attention. Working Paper.
Heidhues, Paul, Kőszegi, Botond, 2008. Competition and price variation when consumers are loss averse. American Economic Review 98 (4), 1245–1268.
Heidhues, Paul, Kőszegi, Botond, 2010. Exploiting naivete about self-control in the credit market. American Economic Review 100 (5), 2279–2303.
Heidhues, Paul, Kőszegi, Botond, 2014. Regular prices and sales. Theoretical Economics 9 (1), 217–251.
Heidhues, Paul, Kőszegi, Botond, 2015. On the welfare costs of naivete in the US credit-card market. Review of Industrial Organization 47 (3), 341–354.
Heidhues, Paul, Kőszegi, Botond, 2017. Naivete-based discrimination. Quarterly Journal of Economics 132 (2), 1019–1054.
Heidhues, Paul, Kőszegi, Botond, Murooka, Takeshi, 2012a. Deception and consumer protection in competitive markets. In: The Pros and Cons of Consumer Protection. Swedish Competition Authority (Chapter 3).

Heidhues, Paul, Kőszegi, Botond, Murooka, Takeshi, 2012b. Inferior Products and Profitable Deception. Working Paper. European School of Management and Technology. https://www.esmt.org/sites/default/files/digital-measures/Inferior%20products-1.pdf.
Heidhues, Paul, Kőszegi, Botond, Murooka, Takeshi, 2016. Exploitative innovation. American Economic Journal: Microeconomics 8 (1), 1–23.
Heidhues, Paul, Kőszegi, Botond, Murooka, Takeshi, 2017. Inferior products and profitable deception. Review of Economic Studies 84 (1), 323–356.
Heidhues, Paul, Johnen, Johannes, Kőszegi, Botond, 2018. Browsing Versus Studying Offers. Working Paper.
Heifetz, Aviad, Shannon, Chris, Spiegel, Yossi, 2007. What to maximize if you must. Journal of Economic Theory 133 (1), 31–57.
Helfrich, Magdalena, Herweg, Fabian, 2017. Salience in Retailing: Vertical Restraints on Internet Sales. Working Paper.
Hendricks, Kenneth, Porter, Robert H., 1988. An empirical study of an auction with asymmetric information. American Economic Review 78 (5), 865–883.
Hendricks, Kenneth, Porter, Robert H., Boudreau, Bryan, 1987. Information, returns, and bidding behavior in OCS auctions: 1954–1969. Journal of Industrial Economics 35 (4), 517–542.
Herweg, Fabian, Mierendorff, Konrad, 2013. Uncertain demand, consumer loss aversion, and flat-rate tariffs. Journal of the European Economic Association 11 (2), 399–432.
Herweg, Fabian, Schmidt, Klaus M., 2015. Loss aversion and inefficient renegotiation. Review of Economic Studies 82 (1), 297–332.
Herweg, Fabian, Müller, Daniel, Weinschenk, Philipp, 2017. Salience in Markets. Working Paper.
Herweg, Fabian, Karle, Heiko, Müller, Daniel, 2018. Incomplete contracting, renegotiation, and expectation-based loss aversion. Journal of Economic Behavior & Organization 145, 176–201.
Herz, Holger, Schunk, Daniel, Zehnder, Christian, 2014. How do judgmental overconfidence and overoptimism shape innovative activity? Games and Economic Behavior 83 (C), 1–23.
Hoffmann, Florian, Inderst, Roman, Ottaviani, Marco, 2014. Persuasion Through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation. Working Paper.
Holmström, Bengt, 1979. Moral hazard and observability. Bell Journal of Economics 10 (1), 74–91.
Holmström, Bengt, 1982. Moral hazard in teams. Bell Journal of Economics 13 (2), 324–340.
Hortaçsu, Ali, Puller, Steven L., 2008. Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market. RAND Journal of Economics 39 (1), 86–114.
Hortaçsu, Ali, Syverson, Chad, 2004. Product differentiation, search costs, and competition in the mutual fund industry: a case study of S&P 500 index funds. The Quarterly Journal of Economics 119 (2), 403–456.
Hortaçsu, Ali, Luco, Fernando, Puller, Steven L., Zhu, Dongni, 2017. Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets. Working Paper.
Hotelling, Harold, 1929. Stability in competition. The Economic Journal 39 (153), 41–57.
Huck, Steffen, Weizsäcker, Georg, 2015. Markets for Leaked Information. Working Paper.
Inderst, Roman, Obradovits, Martin, 2016. Excessive Competition for Headline Prices. Working Paper.
Inderst, Roman, Obradovits, Martin, 2017. Loss Leading with Salient Thinkers. Working Paper.
Inderst, Roman, Ottaviani, Marco, 2012a. Financial advice. Journal of Economic Literature 50 (2), 494–512.
Inderst, Roman, Ottaviani, Marco, 2012b. How (not) to pay for advice: a framework for consumer financial protection. Journal of Financial Economics 105 (2), 393–411.
Jehiel, Philippe, 2005. Analogy-based expectation equilibrium. Journal of Economic Theory 123 (2), 81–104.
Johnen, Johannes, 2017a. Dynamic Competition in Deceptive Markets. Working Paper. ESMT.
Johnen, Johannes, 2017b. Screening Procrastinators with Automatic-Renewal Contracts. Working Paper.

Johnson, Justin P., 2017. Unplanned purchases and retail competition. American Economic Review 107 (3), 931–965.
Kahneman, Daniel, Knetsch, Jack L., Thaler, Richard, 1986. Fairness as a constraint on profit seeking: entitlements in the market. American Economic Review 76 (4), 728–741.
Kaldor, Nicholas, 1950. The economic aspects of advertising. Review of Economic Studies 18 (1), 1–27.
Kamenica, Emir, Mullainathan, Sendhil, Thaler, Richard, 2011. Helping consumers know themselves. American Economic Review 101 (3), 417–422.
Karle, Heiko, Peitz, Martin, 2014. Competition under consumer loss aversion. RAND Journal of Economics 45 (1), 1–31.
Karle, Heiko, Peitz, Martin, 2017. De-targeting: advertising an assortment of products to loss-averse consumers. European Economic Review 95, 103–124.
Karle, Heiko, Schumacher, Heiner, 2017. Advertising and attachment: exploiting loss aversion through pre-purchase information. RAND Journal of Economics 48 (4), 927–948.
Kiss, András, 2014. Salience and Switching. Working Paper. University of Amsterdam.
Klein, Benjamin, Crawford, Robert G., Alchian, Armen A., 1978. Vertical integration, appropriable rents, and the competitive contracting process. Journal of Law and Economics 21 (2), 297–326.
Ko, K. Jeremy, 2012. Disclosure and Price Regulation in a Market with Potentially Shrouded Costs. Mimeo.
Korobkin, Russell, 2003. Bounded rationality, standard form contracts, and unconscionability. University of Chicago Law Review 70 (4), 1203–1295.
Kosfeld, Michael, Schüwer, Ulrich, 2017. Add-on pricing in retail financial markets and the fallacies of consumer education. Review of Finance 21 (3), 1189–1216.
Kőszegi, Botond, 2014. Behavioral contract theory. Journal of Economic Literature 52 (4), 1075–1118.
Kőszegi, Botond, Rabin, Matthew, 2006. A model of reference-dependent preferences. Quarterly Journal of Economics 121 (4), 1133–1166.
Kőszegi, Botond, Rabin, Matthew, 2007. Reference-dependent risk attitudes. American Economic Review 97 (4), 1047–1073.
Kőszegi, Botond, Szeidl, Ádám, 2013. A model of focusing in economic choice. Quarterly Journal of Economics 128 (1), 53–107.
Krähmer, Daniel, 2006. Advertising and conspicuous consumption. Journal of Institutional and Theoretical Economics 162 (4), 661–682.
Kuksov, Dmitri, Wang, Kangkang, 2013. A model of the "it" products in fashion. Marketing Science 32 (1), 51–69.
Kuksov, Dmitri, Xie, Ying, 2012. Competition in a status goods market. Journal of Marketing Research 49 (5), 609–623.
Lacetera, Nicola, Pope, Devin G., Sydnor, Justin, 2010. Heuristic Thinking and Limited Attention in the Car Market. Working Paper. The Wharton School.
Laibson, David I., Yariv, Leeat, 2007. Safety in Markets: An Impossibility Theorem for Dutch Books. Working Paper. Caltech.
Lal, Rajiv, Matutes, Carmen, 1994. Retail pricing and advertising strategies. Journal of Business 67 (3), 345–370.
Landier, Augustin, Thesmar, David, 2009. Financial contracting with optimistic entrepreneurs. Review of Financial Studies 22 (1), 117–150.
Malmendier, Ulrike, Tate, Geoffrey, 2005. CEO overconfidence and corporate investment. Journal of Finance 60 (6), 2661–2700.
Malmendier, Ulrike, Tate, Geoffrey, 2008. Who makes acquisitions? CEO overconfidence and the market's reaction. Journal of Financial Economics 89 (1), 20–43.
Mandler, Michael, 2018. Piracy versus monopoly in the market for conspicuous consumption. Economic Journal 128 (610), 1257–1275.
Manove, Michael, Padilla, A. Jorge, 1999. Banking (conservatively) with optimists. Rand Journal of Economics 30 (2), 324–350.

Martin, Daniel, 2017. Strategic pricing with rational inattention to quality. Games and Economic Behavior 104, 131–145.
Massey, Cade, Thaler, Richard H., 2013. The loser's curse: decision making and market efficiency in the national football league draft. Management Science 59 (7), 1479–1495.
Mathios, Alan D., 2000. The impact of mandatory disclosure laws on product choices: an analysis of the salad dressing market. Journal of Law and Economics 43 (2), 651–677.
Mazali, Rogério, Rodrigues-Neto, José A., 2013. Dress to impress: brands as status symbols. Games and Economic Behavior 82, 103–131.
McMillan, Robert S., 2004. Estimating Demand for Differentiated Products when Choice is Continuous and Complementarities Exist: An Application to the Puzzle of Uniform Pricing. PhD dissertation. Stanford University.
Miao, Chun-Hui, 2010. Consumer myopia, standardization and aftermarket monopolization. European Economic Review 54 (7), 931–946.
Michel, Christian, 2016. Contractual Structures and Consumer Misperceptions—The Case of Warranties. Working Paper.
Michel, Christian, 2017. Market Regulation of Voluntary Add-on Contracts. Working Paper.
Milgrom, Paul R., 1981. Good news and bad news: representation theorems and applications. Bell Journal of Economics 12 (2), 380–391.
Moore, Don A., Healy, Paul J., 2008. The trouble with overconfidence. Psychological Review 115 (2), 502–517.
Moskowitz, Tobias J., Vissing-Jørgensen, Annette, 2002. The returns to entrepreneurial investment: a private equity premium puzzle? American Economic Review 92 (4), 745–778.
Mullainathan, Sendhil, Nöth, Markus, Schoar, Antoinette, 2011. The Market for Financial Advice: An Audit Study. Mimeo. Massachusetts Institute of Technology.
Muris, Timothy, 2002. The Interface of Competition and Consumer Protection. Remarks at Fordham Corporate Law Institute's 29th Annual International Antitrust Law and Policy Conference, New York.
Murooka, Takeshi, 2015. Deception under Competitive Intermediation. Working Paper.
Murooka, Takeshi, Schwarz, Marco A., 2018. The timing of choice-enhancing policies. Journal of Public Economics 157, 27–40.
Nocke, Volker, Peitz, Martin, 2003. Hyperbolic discounting and secondary markets. Games and Economic Behavior 44 (1), 77–97.
Nuñez, Benjamin J.R., 2017. Developing behavioural economics as a practical tool for market authorities. Journal of Antitrust Enforcement 5 (3), 375–406.
Office of Gas and Electricity Markets, 2013. The Retail Market Review—Implementation of Simpler Tariff Choices and Clearer Information. Technical Report.
Pesendorfer, Wolfgang, 1995. Design innovation and fashion cycles. American Economic Review 85 (4), 771–792.
Piccione, Michele, Spiegler, Ran, 2012. Price competition under limited comparability. Quarterly Journal of Economics 127 (1), 135–197.
Rao, Raghunath Singh, Schaefer, Richard, 2013. Conspicuous consumption and dynamic pricing. Marketing Science 32 (5), 786–804.
Rayo, Luis, 2013. Monopolistic signal provision. The B.E. Journal of Theoretical Economics 13 (1), 1–32.
Roll, Richard, 1986. The hubris hypothesis of corporate takeovers. Journal of Business 59, 197–216.
Romer, David, 2006. Do firms maximize? Evidence from professional football. Journal of Political Economy 114 (2), 340–365.
Rosato, Antonio, 2016. Selling substitute goods to loss-averse consumers: limited availability, bargains and rip-offs. RAND Journal of Economics 47 (3), 709–733.
Rotemberg, Julio J., 2005. Customer anger at price increases, changes in the frequency of price adjustment and monetary policy. Journal of Monetary Economics 52 (4), 829–852.

Rotemberg, Julio J., 2010. Altruistic dynamic pricing with customer regret. Scandinavian Journal of Economics 112 (4), 646–672.
Rotemberg, Julio J., 2011. Fair pricing. Journal of the European Economic Association 9 (5), 952–981.
Rothschild, Michael, Stiglitz, Joseph E., 1976. Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. Quarterly Journal of Economics 90 (4), 630–649.
Ru, Hong, Schoar, Antoinette, 2016. Do Credit Card Companies Screen for Behavioral Biases? Working Paper.
Salant, Yuval, Siegel, Ron, 2017. Contracts with Framing. Working Paper.
Salop, Steven C., 1979. Monopolistic competition with outside goods. Bell Journal of Economics 10 (1), 141–156.
Sandroni, Alvaro, Squintani, Francesco, 2007. Overconfidence, insurance, and paternalism. American Economic Review 97 (5), 1994–2004.
Sandroni, Alvaro, Squintani, Francesco, 2010. Overconfidence and Asymmetric Information: The Case of Insurance. Working Paper. University of Warwick.
Schelling, Thomas C., 1960. The Strategy of Conflict. Harvard University Press.
Scherer, Frederic M., Ross, David, 1990. Industrial Market Structure and Economic Performance. Houghton Mifflin Company.
Schumacher, Heiner, 2016. Insurance, self-control, and contract flexibility. European Economic Review 83 (Supplement C), 220–232.
Scitovsky, Tibor, 1950. Ignorance as a source of oligopoly power. American Economic Review 40 (2), 48–53.
Seim, Katja, Vitorino, Maria Ana, Muir, David M., 2016. Drip Pricing When Consumers Have Limited Foresight: Evidence from Driving School Fees. Working Paper.
Shapiro, Carl, 1995. Aftermarkets and consumer welfare: making sense of Kodak. Antitrust Law Journal 63 (2), 483–511.
Shui, Haiyan, Ausubel, Lawrence M., 2004. Time Inconsistency in the Credit Card Market. Available at http://ssrn.com/abstract=586622.
Simon, Herbert A., 1951. A formal theory of the employment relationship. Econometrica 19 (3), 293–305.
Simon, Herbert A., 1955. A behavioral model of rational choice. Quarterly Journal of Economics 69 (1), 99–118.
Simon, Herbert A., 1959. Theories of decision-making in economics and behavioral science. American Economic Review 49 (3), 253–283.
Sims, Christopher A., 2003. Implications of rational inattention. Journal of Monetary Economics 50 (3), 665–690.
Spiegler, Ran, 2006a. Competition over agents with boundedly rational expectations. Theoretical Economics 1 (2), 207–231.
Spiegler, Ran, 2006b. The market for quacks. Review of Economic Studies 73 (4), 1113–1131.
Spiegler, Ran, 2012. Monopoly pricing when consumers are antagonized by unexpected price increases: a "cover version" of the Heidhues–Koszegi–Rabin model. Economic Theory 51, 695–711.
Spiegler, Ran, 2014. Competitive framing. American Economic Journal: Microeconomics 6 (3), 35–58.
Spiegler, Ran, 2015. On the equilibrium effects of nudging. Journal of Legal Studies 44 (2), 389–416.
Spiegler, Ran, 2016. Choice complexity and market competition. Annual Review of Economics 8, 1–25.
Spinnewijn, Johannes, 2013. Insurance and perceptions: how to screen optimists and pessimists. Economic Journal 123 (569), 606–633.
Syverson, Chad, 2011. What determines productivity? Journal of Economic Literature 49 (2), 326–365.
Thaler, Richard H., 1988. Anomalies: the winner's curse. Journal of Economic Perspectives 2 (1), 191–202.
Thaler, Richard H., Benartzi, Shlomo, 2004. Save more tomorrow: using behavioral economics to increase employee saving. Journal of Political Economy 112 (1), S164–S187.
Thaler, Richard H., Sunstein, Cass R., 2003. Libertarian paternalism. American Economic Review 93 (2), 175–179.

Thaler, Richard H., Sunstein, Cass R., 2008. Nudge: Improving Decisions About Health, Wealth, and Happiness. Penguin Books.
Tirole, Jean, 1988. The Theory of Industrial Organization. MIT Press.
Varian, Hal R., 1980. A model of sales. American Economic Review 70 (4), 651–659.
Vickers, John, 1985. Delegation and the theory of the firm. The Economic Journal 95, 138–147.
Warren, Patrick L., Wood, Daniel H., 2014. The political economy of regulation in markets with naïve consumers. Journal of the European Economic Association 12 (6), 1617–1642.
Wenzel, Tobias, 2014. Consumer myopia, competition and the incentives to unshroud add-on information. Journal of Economic Behavior & Organization 98 (C), 89–96.
Williamson, Oliver E., 1971. The vertical integration of production: market failure considerations. American Economic Review 61 (2), 112–123.
Wilson, Chris M., 2010. Ordered search and equilibrium obfuscation. International Journal of Industrial Organization 28 (5), 496–506.
Yoganarasimhan, Hema, 2012. Cloak or flaunt? The fashion dilemma. Marketing Science 31 (1), 74–95.
Zenger, Hans, 2013. Why firms' exploitation of consumer myopia may benefit myopic consumers. Economics Letters 118 (2), 307–309.
Zhou, Jidong, 2011. Reference dependence and market competition. Journal of Economics and Management Strategy 20 (4), 1073–1097.