NAÏVETÉ-BASED DISCRIMINATION[*]

PAUL HEIDHUES AND BOTOND KŐSZEGI

We initiate the study of naïveté-based discrimination, the practice of conditioning offers on external information about consumers' naïveté. Knowing that a consumer is naive increases a monopolistic or competitive firm's willingness to generate inefficiency to exploit the consumer's mistakes, so naïveté-based discrimination is not Pareto-improving, can be Pareto-damaging, and often lowers total welfare when classical preference-based discrimination does not. Moreover, the effect on total welfare depends on a hitherto unemphasized market feature: the extent to which the exploitation of naive consumers distorts trade with different types of consumers. If the distortion is homogeneous across naive and sophisticated consumers, then under an arguably weak and empirically testable condition, naïveté-based discrimination lowers total welfare. In contrast, if the distortion arises only for trades with sophisticated consumers, then perfect naïveté-based discrimination maximizes social welfare, although imperfect discrimination often lowers welfare. If the distortion arises only for trades with naive consumers, then naïveté-based discrimination has no effect on welfare. We identify applications for each of these cases. In our primary example, a credit market with present-biased borrowers, firms lend more than is socially optimal to increase the amount of interest naive borrowers unexpectedly pay, creating a homogeneous distortion. The condition for naïveté-based discrimination to lower welfare is then weaker than prudence. *JEL Codes:* D21, D49, D69, L19.

I. INTRODUCTION

What is the welfare effect of firms knowing more about consumers? The classical approach to answering this old question

[*]We thank Dan Benjamin, Michael Grubb, Fabian Herweg, Takeshi Murooka, Christopher Snyder, Rani Spiegler, the editor, anonymous referees, and numerous seminar and conference audiences for helpful comments. Financial support from the European Research Council (Starting Grant #313341) is also gratefully acknowledged. Part of the research was carried out while Heidhues was visiting the Institute for Advanced Studies at CEU and part while he was visiting Dartmouth, and he is grateful for both institutions' hospitality. Another part was carried out while both authors were visiting WZB, and we are grateful for their hospitality as well. Finally, we would like to thank Johannes Johnen and Mirjam Szilléry for research assistance.

The Quarterly Journal of Economics (2017), 1019–1054. doi:10.1093/qje/qjw042.
Advance Access publication on November 15, 2016.

in economics, in the vast literatures on price discrimination and privacy, typically presumes that the information firms acquire about consumers pertains to preferences.[1] A growing literature in behavioral economics, however, documents that many consumers are naive about key fees they will incur, raising the possibility that some of the information firms acquire about consumers pertains to naïveté—and hence firms engage in *naïveté-based discrimination*. This possibility is especially relevant since firms have both the incentive and—with the ever-increasing amount of information they collect—the ability to differentially target profitable naive consumers.

In this article, we initiate the study of naïveté-based discrimination, identifying its welfare effects in key canonical environments. We begin in Section II by contrasting the logic of preference-based discrimination and naïveté-based discrimination in a simple example. Suppose that a monopolist bank-account provider faces a population of consumers, some of whom overdraft rarely, and some of whom overdraft often. To serve everyone as well as make extra money on consumers who overdraft often, the bank charges a low account maintenance fee and a high overdraft fee, foreclosing some socially beneficial overdraft transactions for all consumers. What happens, then, if the bank knows which consumers will overdraft more and can offer different fees to them?

The answer depends crucially on why these consumers overdraft more. If they do so because they value the convenience more highly—a preference-based explanation—then the bank lowers the overdraft fee on them and instead efficiently extracts surplus through the account maintenance fee, increasing social welfare. While in this case discrimination lowers the welfare of high-value consumers, in notable other cases—such as when low-value consumers are initially excluded—discrimination is Pareto-improving. Importantly, these effects rely on market power: if the bank operates in a perfectly competitive environment, then with or without discrimination marginal cost pricing obtains, so discrimination has no effect on welfare.

1. The literature distinguishes between three main types of price discrimination. Under first-degree or perfect price discrimination, a firm knows consumers' preferences perfectly. Under third-degree price discrimination, a firm can offer different deals to different groups of consumers based on some observable characteristic. First- and third-degree discrimination rely on external information about consumers, which is also our main interest in this article. In contrast, under second-degree price discrimination a firm gives different deals to different consumers by inducing self-selection among them.

In contrast, if consumers who overdraft more do so unexpectedly and by mistake—a naïveté-based explanation—then they are not willing to pay a lot for account maintenance, so the only way to extract profits from them is through the overdraft fee. If the bank knows who will overdraft more, it, therefore, raises the overdraft fee on these consumers, lowering consumer and social welfare. In addition, a bank can profit from consumers' mistakes for any level of competition, so naïveté-based discrimination remains welfare-relevant even in perfectly competitive environments. In fact, in a competitive market naïveté-based discrimination may be Pareto-damaging: it can hurt both naive and sophisticated consumers without affecting firms' profits.

In the rest of the article, we systematically characterize the welfare effects of naïveté-based discrimination, defined as discrimination using any information about the likelihood that a consumer is naive. In Section III, we analyze a simple reduced-form pricing model in the tradition of Gabaix and Laibson (2006) where firms can, potentially inefficiently, induce naive consumers to pay unexpected charges. For monopolistic as well as competitive markets, we confirm a basic implication illustrated by our example: that unlike preference-based discrimination, naïveté-based discrimination is never Pareto-improving. In addition, we show that the aggregate welfare implications depend qualitatively on a hitherto unemphasized market feature: the extent to which the exploitation of naive consumers distorts trade with sophisticated versus naive consumers. Here, we illustrate these general insights using specific applications we develop in Section IV.

In Section IV.A, we analyze a model of the credit market with present-biased consumers based on Heidhues and Kőszegi (2010). Consumers choose between contracts specifying a loan amount and an interest rate, and later decide how much of their loan to repay early. In equilibrium, naive but not sophisticated consumers pay unexpected interest on their loan, and to increase such interest payments, lenders extend more credit than is socially optimal. Because naive and sophisticated consumers getting the same deal overborrow by the same amount, this market features a "homogeneous distortion." Naïveté-based discrimination lowers welfare if consumers' consumption-utility function satisfies prudence, which is a widely made and empirically supported assumption on the utility function. More generally, with homogeneous distortions naïveté-based discrimination lowers welfare if an arguably weak and empirically testable condition on the distortion holds. Intuitively, discrimination increases overlending to consumers more likely to be naive, and decreases overlending to consumers more likely to be sophisticated. Because an increase in a preexisting distortion is more costly than an identical decrease is beneficial, the net effect is often negative.

In Section IV.B, we develop a model of markets with expensive add-ons inspired by Gabaix and Laibson (2006) and especially Grubb (2015), framing the model in the context of mobile phones. Providers with zero marginal cost of calls offer plans consisting of a monthly fee and an amount of free airtime, charging a high price for additional minutes. Each consumer expects to forgo her lowest-value minutes of calling to avoid overage charges, but naive consumers fail to do so. Because only sophisticated consumers undertake socially inefficient avoidance, this market features a "sophisticated-side distortion." In such settings, perfect naïveté-based discrimination always maximizes welfare. Intuitively, if a firm knows that a consumer is sophisticated and hence cannot be exploited, then it does not charge an overage fee, so no distortion arises; and if the firm knows that the consumer is naive, then it can charge a high overage fee without triggering an inefficiency.

Finally, in Section IV.C we present a model in which retailers can offer useless but costly extras to a purchased product. Consumers expect to buy a basic car and hence choose a car seller with the lowest base price. Once a consumer is at the shop, the seller can offer extras—such as rust proofing or paint protection—that create no value, which naive but not sophisticated consumers accept. Since the socially wasteful cost of providing the extras arises only for naive consumers, this market generates a "naive-side distortion." In such an environment, naïveté-based discrimination has no impact on welfare. Intuitively, because the extras do not affect the profits from sophisticated consumers, the seller maximizes the ex post profits from naive consumers, leading her to offer the same extras independently of what she knows about consumers' naïveté.

In Section V, we discuss some important empirical issues raised by our model. For situations of model uncertainty, we suggest ways to use market data to determine which trades are distorted by the exploitation of naive consumers. We also provide an empirical method to assess whether the condition for naïveté-based discrimination to lower welfare with homogeneous distortions holds.

In Section VI, we relate our article to the empirical and theoretical literatures on markets with naive consumers and on price discrimination and privacy. Although direct evidence that firms engage in naïveté-based discrimination is limited, we argue based on theoretical and empirical considerations that they likely do.

While this article covers many of the relevant cases, it is just a first step in understanding the welfare effects of naïveté-based discrimination. In Section VII, we point out some of the main issues that call for further research. We emphasize that homogeneous, sophisticated-side, and naive-side distortions may be present in a market at the same time and interact with each other as well as with classical preference-based discrimination and naïveté-based screening, and we do not know the consequences of these interactions.

II. Example

In this section, we present a simple example to contrast preference-based and naïveté-based discrimination in the same setting.[2] Suppose that a monopolist offers bank accounts with add-on overdraft protection, choosing an account maintenance fee $\tilde{f}$ as well as an overdraft fee per dollar of overdraft $\tilde{a} \leqslant 2$. The bank's marginal cost of providing an account, including overdraft protection, is zero. All consumers value account maintenance without overdraft protection at \$3. But consumers differ in their demand for overdrafts: the demand curve of low-demand consumers is $2 - \tilde{a}$, whereas that of high-value consumers is $4 - \tilde{a}$.

High-demand consumers' extra \$2 of overdraft usage is consistent with both preferences and unanticipated mistakes. Under the former, classic interpretation, the two demand curves reflect low-value and high-value consumers' surpluses from the convenience of overdrafting, respectively.[3] Under the latter interpretation, both sophisticated and naive consumers correctly

2. The example is a numerical special case of the model of add-ons we study in Section IV.A, which yields a homogeneous distortion.

3. Denoting overdraft usage by o, a utility function for bank accounts consistent with this interpretation is $5 - \frac{(2-o)^2}{2} - \tilde{a}o - \tilde{f}$ for low-value consumers and $11 - \frac{(4-o)^2}{2} - \tilde{a}o - \tilde{f}$ for high-value consumers. First, when $o = 0$, both consumers get a utility of $3 - \tilde{f}$, so they are willing to pay \$3 for an account without overdraft protection. Second, differentiating the utility functions with respect to o yields the first-order conditions $2 - o - \tilde{a} = 0$ and $4 - o - \tilde{a} = 0$ for low-value and high-value consumers, respectively. Rearranging gives the specified demand curves.

understand that their surplus from overdrafts is captured by the low-value demand curve $2 - \tilde{a}$, and they expect to overdraft by exactly $2 - \tilde{a}$. But naive consumers incur an additional \$2 of overdrafts that do not generate any value to them, for instance, by mistiming some bill payments from the account.[4]

Our main interest is in identifying what happens if the bank can distinguish low-demand and high-demand consumers and offer different prices to them. Denoting the share of high-demand consumers in a group getting the same offer by α, such perfect discrimination corresponds to an increase in α to 1 for high-demand consumers, and a decrease in α to 0 for low-demand consumers.

The effect of discrimination depends crucially on whether the preference-based or the naïveté-based model applies. We start with the preference-based model, first deriving outcomes for a fixed α. For any overdraft fee $\tilde{a}$, the bank either (i) chooses the maintenance fee $\tilde{f}$ so that low-value consumers are indifferent between accepting and not accepting—serving both types and giving high-value consumers rent—or (ii) chooses the maintenance fee $\tilde{f}$ so that high-value consumers are indifferent between accepting and not accepting—excluding low-value consumers and giving high-value consumers no rent. Standard arguments imply that the optimal strategy (i) has $\tilde{a} = 2\alpha$ and dominates for $\alpha < \frac{1}{2}$, whereas the optimal strategy (ii) has $\tilde{a} = 0$ and dominates for $\alpha > \frac{1}{2}$.[5]

Suppose that low-value consumers are in the majority in the population ($\alpha < \frac{1}{2}$), which implies that without discrimination the bank uses an inefficient, positive overdraft fee. Perfect discrimination leads the bank to lower the overdraft fee to 0 for all consumers, increasing social welfare to first-best. Intuitively, to serve the majority low-value consumers without discrimination,

4. In this interpretation, both sophisticated and naive consumers believe that their utility from a bank account is the same as that of a low-value consumer above, which by note 3 is $5 - \frac{\tilde{a}^2}{2} - \tilde{a}(2 - \tilde{a}) - \tilde{f}$. Sophisticated consumers are correct in this belief. But a naive consumer pays $2\tilde{a}$ more in overdraft fees than she expects, so her actual utility is $5 - \frac{\tilde{a}^2}{2} - \tilde{a}(4 - \tilde{a}) - \tilde{f}$.

5. Under strategy (i), the bank chooses $\tilde{f} = 5 - \frac{\tilde{a}^2}{2} - \tilde{a}(2 - \tilde{a})$, earning $5 - \frac{\tilde{a}^2}{2} - \tilde{a}(2 - \tilde{a}) + (1 - \alpha)(2 - \tilde{a})\tilde{a} + \alpha(4 - \tilde{a})\tilde{a} = 5 - \frac{\tilde{a}^2}{2} + 2\alpha\tilde{a}$ for any $\tilde{a} \leqslant 2$. The first-order condition with respect to $\tilde{a}$ gives $\tilde{a} = 2\alpha$. Hence, strategy (i) earns $5 + 2\alpha^2$. Under strategy (ii), because the monopolist extracts all of high-value consumers' surplus, it sets $\tilde{a} = 0$ to maximize that surplus, and then charges $\tilde{f} = 3 + (\frac{1}{2})16 = 11$, earning 11α. Furthermore, $5 + 2\alpha^2 > 11\alpha$—and hence strategy (i) dominates—if $\alpha < \frac{1}{2}$, while $5 + 2\alpha^2 < 11\alpha$—and hence strategy (ii) dominates—if $\alpha > \frac{1}{2}$.

the bank—unable to make a separate offer—must give away rent to high-value consumers. To lower that rent, the bank sets an inefficiently high overdraft fee. Under perfect discrimination, however, the bank knows which consumers are high-value and can make them a separate offer. It can therefore extract all rent from high-value consumers through a high maintenance fee, and there is no need for an inefficiently high overdraft fee. While discrimination lowers the utility of high-value consumers, consumers who have such a high willingness to pay for bank accounts may on average be relatively wealthy, so this distributional impact may not be too adverse.

In addition, if high-value consumers are in the majority in the population ($\alpha > \frac{1}{2}$), then discrimination is Pareto-improving. In this case, without discrimination the bank excludes low-value consumers to extract all rent from the majority high-value consumers, so with or without discrimination all consumers receive zero utility. But perfect discrimination allows the bank to serve low-value consumers with a different account, increasing profits. And imperfect discrimination can benefit consumers as well as the bank: if the bank sorts a minority of high-value consumers into the low-value pool (yielding $0 < \alpha < \frac{1}{2}$), then it adopts strategy (i) for this mixed pool, giving rent to high-value consumers in the pool.

Finally, in a perfectly competitive market consumers always receive the efficient overdraft fee $\tilde{a} = 0$ and also pay the competitive maintenance fee $\tilde{f} = 0$, so discrimination has no effect on individual or total welfare. Intuitively, if a bank sold an account with a positive overdraft fee, then a competitor could profitably attract consumers by offering free overdrafts and capturing most of the increase in consumer and social surplus through an increase in the maintenance fee.[6]

We now consider the naïveté-based model. Paralleling our analysis above, we first derive outcomes for a fixed share α of naive consumers. Note that sophisticated consumers are identical to the low-value consumers above. While naive consumers'

6. Formally, suppose for the sake of contradiction that some consumers accept an account with $\tilde{a} > 0$. Since these consumers overdraft inefficiently little, consumer welfare (which in a competitive market equals total welfare) is not maximized subject to firms making zero profits. Now the prices $\tilde{f} = \tilde{a} = 0$ maximize both types' welfare subject to zero profits, so at least one type receives strictly lower welfare than with $\tilde{f} = \tilde{a} = 0$. Hence, for a sufficiently small $\epsilon > 0$ the prices $\tilde{f} = \epsilon, \tilde{a} = 0$ attract some consumers and—in contradiction to a perfectly competitive market—earn positive profits.

demand for overdrafts is identical to that of high-value consumers above, since their extra overdraft usage reflects unanticipated mistakes, their willingness to pay for a bank account equals that of sophisticated consumers. These considerations mean that pricing strategy (ii) above—setting the maintenance fee to make high-value consumers indifferent—would lead to zero demand and zero profits. In contrast, strategy (i)—setting the maintenance fee to make low-value consumers indifferent—leads all consumers to accept and is exactly as profitable as before. Hence, the bank always uses the optimal strategy (i), charging $\tilde{a} = 2\alpha$. Sophisticated consumers' utility is zero, and since naive consumers make an extra \$2 of overdrafts at rate 2α, their utility is -4α.

The foregoing implies that for any share of naive consumers, perfect naïveté-based discrimination leads the bank to increase the overdraft fee, lowering naive consumers' welfare. Intuitively, the bank faces a trade-off: a higher overdraft fee raises the amount naive consumers unexpectedly pay ex post, but it induces more inefficiency that all consumers must be compensated for ex ante. If the bank knows that a consumer is naive and hence the ex post profits are forthcoming with certainty, then it is willing to generate more inefficiency to exploit the consumer's mistake. Crucially, because a naive consumer does not appreciate how much she will overdraft, setting the overdraft fee to zero and charging for all of her overdraft demand efficiently through the maintenance fee is not a viable pricing strategy.

The harm on the naive side of the market implies that naïveté-based discrimination is never Pareto-improving, and drives another sharp contrast with preference-based discrimination: whereas perfect preference-based discrimination maximizes total welfare, among different levels of naïveté-based discrimination perfect discrimination minimizes total welfare.[7] Worse,

7. We show that starting from any share $\alpha \in (0, 1)$ of naive consumers, perfect discrimination strictly lowers total welfare; this implies that perfect discrimination leads to lower welfare than any imperfect discrimination. When selling to a naive consumer, a bank earns $\tilde{f} + \tilde{a}(4 - \tilde{a})$ and (by note 4) the consumer receives utility $5 - \frac{\tilde{a}^2}{2} - \tilde{a}(4 - \tilde{a}) - \tilde{f}$, so that total welfare is $5 - \frac{\tilde{a}^2}{2}$. Similarly, when selling to a sophisticated consumer the bank earns $\tilde{f} + \tilde{a}(2 - \tilde{a})$ and the consumer receives utility $5 - \frac{\tilde{a}^2}{2} - \tilde{a}(2 - \tilde{a}) - \tilde{f}$, so that total welfare is again $5 - \frac{\tilde{a}^2}{2}$. Without discrimination, the bank sets an overdraft fee $\tilde{a} = 2\alpha$ for all consumers, yielding a total welfare of $5 - 2\alpha^2$. With discrimination, the bank sets $\tilde{a} = 0$ for a share $1 - \alpha$ of consumers and $\tilde{a} = 2$ for a share α, yielding a total welfare of $5 - 2\alpha < 5 - 2\alpha^2$.

because perfect naïveté-based discrimination leaves the welfare of sophisticated consumers unchanged (at the outside option of zero) and raises the bank's profit, naive consumers bear over 100% of the social welfare loss. In as much as consumers who make financial mistakes have lower incomes than average (Calvet, Campbell, and Sodini 2007), such a distributional impact is extremely adverse.

Finally, while preference-based discrimination is only relevant with market power, the impact of naïveté-based discrimination on total welfare is independent of the level of competition.[8] Intuitively, the transfer generated by a naive consumer's mistake is fully unappreciated by the consumer, so it is not subject to competition at all. As a result, competition leaves the overdraft fee, and hence also total welfare, unchanged. But because competition affects the maintenance fee, it affects the distribution of welfare. In particular, the prospect of profits from naive consumers' overdraft expenditures leads banks to compete more aggressively on the maintenance fee, benefiting sophisticated consumers who get the same account. This has an important implication for discrimination: under perfect competition, perfect naïveté-based discrimination lowers the welfare of both naive and sophisticated consumers.[9] Intuitively, naive consumers are hurt because banks exploit their mistake more aggressively; sophisticated consumers are hurt because (getting a separate offer) they do not receive any of the profits banks make on naive consumers. Hence, again unlike

8. To see this, we think of a bank's problem in two parts: (i) choosing the optimal level of utility $\hat{u}$ to give to sophisticated consumers; and (ii) choosing the optimal prices that give sophisticated consumers utility $\hat{u}$. The level of competition affects consumers' alternative options, so it affects the optimal $\hat{u}$ in (i). But we establish that the optimal overdraft fee in (ii) is independent of $\hat{u}$. Equating $\hat{u}$ with the indirect utility of sophisticated consumers in note 4 yields $\tilde{f} = 5 - \frac{\tilde{a}^2}{2} - \tilde{a}(2 - \tilde{a}) - \hat{u}$. Substituting this into the bank's per consumer profit function $\tilde{f} + (1 - \alpha)(2 - \tilde{a})\tilde{a} + \alpha(4 - \tilde{a})\tilde{a}$ and maximizing with respect to $\tilde{a}$ establishes that $\tilde{a} = 2\alpha$ for any $\hat{u}$.

9. The bank's zero-profit condition is $\tilde{f} + (1 - \alpha)(2 - \tilde{a})\tilde{a} + \alpha(4 - \tilde{a})\tilde{a} = 0$. Using that (by the preceding note) $\tilde{a} = 2\alpha$, this yields $\tilde{f} = -4\alpha$. Plugging into the indirect utilities from note 4 gives that sophisticated consumers' utility is $5 + 2\alpha^2$, and naive consumers' utility is $5 + 2\alpha^2 - 4\alpha$. Perfect naïveté-based discrimination lowers the utility of sophisticated consumers since $5 < 5 + 2\alpha^2$ for any $\alpha \in (0, 1)$. And perfect naïveté-based discrimination lowers the utility of naive consumers since $3 < 5 + 2\alpha^2 - 4\alpha$, or $0 < 2(1 - \alpha)^2$, for any $\alpha \in (0, 1)$.

basic specifications of preference-based discrimination, naïveté-based discrimination can be Pareto-damaging.[10]

III. A Reduced-Form Model of Naïveté-Based Discrimination

In this section, we employ a reduced-form model to study the welfare effects of naïveté-based discrimination in different economic environments. Closely following the spirit of Gabaix and Laibson's (2006) distinction between a base-good price and a shrouded price, we distinguish between an "anticipated price" that all consumers understand and an "additional price" that naive consumers ignore. We use the different terminology to highlight that a naive consumer may anticipate some charges beyond the base-good price, and that she may underestimate spending, not because prices are hidden but because she mispredicts her own behavior. We also allow for more types of distortions from the additional price than do Gabaix and Laibson.

III.A. Setup

We use a Hotelling-type model of pricing with horizontally differentiated products. Consumers with tastes y distributed uniformly on [0, 1] are interested in buying at most one unit of one product, for which they have gross value v. Each consumer is, independently of her taste, naive with probability α and sophisticated with probability $1 - \alpha$. Two firms with identical marginal costs of production c offer products located at $l = 0$ and $l = 1$, respectively. The firms simultaneously choose anticipated prices $f_l \in \mathbb{R}$ and additional prices $a_l \in [\overline{a}, a_{max}]$, where $a_{max} > \overline{a} \geqslant 0$.[11]

10. Our analysis assumes that the bank can offer only one account to a pool of consumers. Under the preference-based model, a monopolist bank may be able to increase profits by offering a second package that stipulates a discount on overdrafts in exchange for a higher maintenance fee, inducing high-value consumers to voluntarily separate. Such screening does not affect any of our qualitative conclusions regarding the effects of discrimination on individual and social welfare. Importantly, screening is impossible in the naïveté-based model, since in that model consumers have the same beliefs regarding their demand for overdrafts. This observation yields an interesting conclusion: a pattern where there is heterogeneity in overdraft demand yet we do not see the bank advertising quantity discounts on overdrafts is more consistent with the naïveté-based model than with the preference-based model.

11. We think of $\overline{a}$ as the highest additional price a firm can impose without creating a distortion. Our specification makes the assumption—for equilibrium payoffs irrelevant—that a firm never chooses a lower additional price.

A naive consumer does not take the additional prices into account when making purchase decisions; but if she buys product l, she ends up paying a_l as well. A sophisticated consumer anticipates the additional price and avoids paying it. In addition, both types incur 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's outside option has gross utility 0, but it is available only at the endpoints of [0, 1], so it has utility $-t \min\{y, 1 - y\}$.[12]

Crucially, we posit that the additional price a_l creates a "distortionary impact" $k(a_l)$ that adds to the social cost of some or all trades. The function $k(\cdot)$ is three times continuously differentiable, with $k(\overline{a}) = k'(\overline{a}) = 0$, $k''(a_l) > 0$ for all a_l, and $k'(a_{max}) \geqslant 1$. We distinguish three conceptually and economically relevant cases based on the trades affected by the distortionary impact:

i. *Homogeneous distortions.* The cost $k(a_l)$ falls on trades with both naive and sophisticated consumers: all consumers expect their utility from purchasing product l to be $v - f_l - t|y - l|$, naive consumers' utility is actually $v - f_l - a_l - t|y - l|$, and firm l's cost of serving a consumer of either type is $c + k(a_l)$.[13]
ii. *Sophisticated-side distortions.* The cost $k(a_l)$ falls only on trades with sophisticated consumers: all consumers expect their utility from purchasing product l to be $v - f_l - k(a_l) - t|y - l|$, a naive consumer's utility is actually $v - f_l - a_l - t|y - l|$, and firm l's cost of serving a consumer of either type is c.[14]
iii. *Naive-side distortions.* The cost $k(a_l)$ falls only on trades with naive consumers: all consumers expect their utility from purchasing product l to be $v - f_l - t|y - l|$, a naive consumer's utility is actually $v - f_l - a_l - t|y - l|$, firm

12. This formulation of the outside option was first introduced by Benabou and Tirole (2016). In the classical Hotelling model—where the utility from the outside option is fixed—t affects both the level of competition and the attractiveness of the product relative to the outside option. Benabou and Tirole's formulation abstracts away from the second effect and hence is more appropriate for studying the pure effect of competition on outcomes.

13. The above formulation assumes that $k(a_l)$ is borne by the firm rather than the consumer. Our analysis will establish that this makes no difference.

14. Unlike with homogeneous distortions, it now matters whether the consumer or the firm pays $k(a_l)$. Because this seems empirically more relevant in our applications, we assume that the consumer pays.

l's cost of serving a sophisticated consumer is c, and its cost from serving a naive consumer is $c + k(a_l)$. We also consider situations in which the consumer bears the distortionary impact. Then, naive consumers' utility is $v - f_l - a_l - k(a_l) - t|y - l|$, and the firm's cost of serving a consumer of either type is c.

Throughout, we solve for symmetric pure-strategy Nash equilibria of the game played between firms, assuming that firms correctly predict consumers' behavior. We suppose that $v > c$, and that all consumers purchase when indifferent. We define social welfare as the sum of firms' profits and the population-weighted sum of naive and sophisticated consumers' utilities.

Our main interest is in the welfare effects of naïveté-based discrimination, which we think of as using information about α to design offers. Without discrimination, firms face one pool of consumers with $\alpha = \alpha_{ns}$ and make a single offer to all consumers. With discrimination, firms sort consumers in an identical way into two pools with $\alpha = \alpha_n > \alpha_{ns}$ and $\alpha = \alpha_s < \alpha_{ns}$, respectively, and can make different offers to the two pools.[15] A theoretically relevant extreme case is perfect naïveté-based discrimination, where $\alpha_n = 1$ and $\alpha_s = 0$.

We discuss some important issues related to the simple setup. First, as is common in the literature, we have made modeling choices to avoid having to study second-degree price discrimination (i.e., screening) simultaneously with third-degree price discrimination. Specifically, we assume that (although they behave differently ex post) naive and sophisticated consumers have the same beliefs at the time of purchase, and by implication firms know consumers' beliefs. Previous work (Eliaz and Spiegler 2006, 2008, for instance) has already shown that if firms do not know consumers' beliefs, then they often screen consumers according to beliefs. But because consumers with the same beliefs (and preferences) choose from any menu in the same way, it is impossible to screen them. The implicit assumption behind our approach is that consumers have already been screened into different pools according to beliefs, and we are studying naïveté-based discrimination within each pool—without studying the screening problem

15. The welfare effects are the same if firms receive a single binary signal on the aggregate share of naive consumers, for instance by conducting research on how naive the population at large is.

in the background or the interaction of the two problems.[16] As a partial justification of this approach in a specific setting, in Online Appendix A we consider a model of a credit market in which firms do not know beliefs, and show that the welfare implications of naïveté-based discrimination are the same as when firms know beliefs.

Second, our framework assumes that any single market features only one type of distortion (homogeneous, sophisticated-side, or naive-side). In many real-life settings, however, multiple types of distortions may be present at the same time. In insurance contracting, for example, a naive-side distortion arises if naive but not sophisticated consumers are left unexpectedly underinsured; and a sophisticated-side distortion also arises if sophisticated but not naive consumers exert costly effort to claim an expected benefit. In Online Appendix B, we briefly consider a market with multiple types of distortions, but leave a full analysis of such combined cases for future work.

Third, although our reduced-form model takes the type of distortion as exogenous, in our applications we derive this type—as well as the functional form of $k(\cdot)$—endogenously from the economic fundamentals of the market. A few general principles emerge:

i. If the distortionary impact arises from ex ante decisions—the purchase itself or choices close to it—then the distortion is homogeneous. Since naive and sophisticated consumers have the same ex ante beliefs and preferences and hence make the same ex ante choices, any distortion that is generated at the ex ante stage must apply equally to them. More generally, the distortion is homogeneous whenever—as in our banking example in Section II—the additional price that naive but not sophisticated consumers pay is nondistortionary, and consumers are otherwise identical. Since naive consumers have the same preferences and make the same distortionary choices

16. The vast majority of research on classical third-degree price discrimination also makes implicit or explicit assumptions to rule out (nontrivial) screening issues. Typically, researchers assume either that a firm can only choose a linear price, or that each individual consumer has unit demand. We are aware of only two papers that study third-degree price discrimination when screening is also going on, Herweg and Müller (2014) and Bergemann, Brooks, and Morris (2015).

as sophisticated consumers, they generate the same distortion.

ii. A sophisticated-side distortion arises if there is a costly behavior aimed at avoiding a fee—or at claiming a rebate or other incentive—that sophisticated consumers undertake but naive consumers do not.

iii. Any distortion associated with the actual payment of the additional price is a naive-side distortion. Collecting the additional price can impose costs on the firm, and facing unexpected expenditures can distort a naive person's consumption across states, products, or time periods.

In some environments, there is nontrivial uncertainty as to which model of naïveté is appropriate, and hence which of our cases applies. Then, an environment-specific empirical test may be necessary. We give examples of such tests in the context of add-on pricing in Section V.

III.B. Homogeneous Distortions

We begin our analysis with homogeneous distortions, where the distortionary impact arises on both sides of the market. To analyze the welfare effect of information about α, we solve for the properties of equilibrium given α. Suppose that in equilibrium, firm l provides a perceived utility gross of transportation costs of $\hat{u}_l$ to consumers. Then, its prices f_l, a_l must solve

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

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

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

Since the equilibrium is symmetric, all consumers buy from the closest firm. Hence, the deadweight loss (DWL) relative to first-best—where consumers buy from the closest firm and $a_0 = a_1 = \overline{a}$—is

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

The foregoing considerations imply that naïveté-based discrimination strictly lowers welfare if

$$DWL(\alpha_{ns}) < \mu_n DWL(\alpha_n) + \mu_s DWL(\alpha_s),$$

where μ_n and μ_s are the population shares of the more naive and more sophisticated pools, respectively. Since we must have $\alpha_{ns} = \mu_n\alpha_n + \mu_s\alpha_s$, the necessary and sufficient condition for naïveté-based discrimination to strictly lower welfare for any α_{ns}, α_n, α_s is that $DWL(\alpha)$ is strictly convex on (0, 1). The first derivative of $DWL(\alpha)$ equals

$$DWL'(\alpha) = k'(a(\alpha))a'(\alpha) = \frac{k'(a(\alpha))}{k''(a(\alpha))},$$

where the second equality follows from totally differentiating equation (1). Using that $a(\alpha)$ is strictly increasing in α, $DWL(\alpha)$ is strictly convex if and only if $\frac{k'(a)}{k''(a)}$ is strictly increasing in the relevant range. Hence:

PROPOSITION 1 (HOMOGENEOUS DISTORTIONS). Given α, there is a unique symmetric pure-strategy Nash equilibrium, which has additional price $a(\alpha) = (k')^{-1}(\alpha)$. For any $\underline{\alpha}$ and $\overline{\alpha} > \underline{\alpha}$, the following are equivalent:

i. $\frac{k'(a)}{k''(a)}$ is strictly increasing in a over the interval $[a(\underline{\alpha}), a(\overline{\alpha})]$.
ii. For any $\alpha_{ns}, \alpha_n, \alpha_s \in [\underline{\alpha}, \overline{\alpha}]$, naïveté-based discrimination strictly lowers social welfare.

Conversely, the following are equivalent:

i′. $\frac{k'(a)}{k''(a)}$ is strictly decreasing in a over the interval $[a(\underline{\alpha}), a(\overline{\alpha})]$.
ii′. For any $\alpha_{ns}, \alpha_n, \alpha_s \in [\underline{\alpha}, \overline{\alpha}]$, naïveté-based discrimination strictly increases social welfare.

Naïveté-based discrimination leads to an increase in the additional price for the more naive pool and to a decrease in the additional price for the more sophisticated pool. Because an increase in a preexisting distortion is more costly than an identical decrease is beneficial, the net effect is often negative. Reversing, mitigating, or exacerbating this tendency is that the above effects on the additional prices may be asymmetric. In particular, if the additional price decreases for the sophisticated pool sufficiently more than it increases for the naive pool, then naïveté-based discrimination raises welfare. This would, however, require that the marginal distortionary impact increases much faster for increases in a than it decreases for decreases in a. The condition in Proposition 1 that $\frac{k'(a)}{k''(a)}$ is increasing in a, a property called decreasing absolute convexity, rules this out, providing a necessary and sufficient condition for any naïveté-based discrimination to be welfare-decreasing. Conversely, the necessary and sufficient condition for any naïveté-based discrimination to be welfare-increasing is that $\frac{k'(a)}{k''(a)}$ is decreasing in a.

It is worth emphasizing that the level of market power (t) does not affect the welfare distortion generated in equilibrium. From the perspective of an individual firm, competition affects consumers' outside option and hence the gross perceived utility ($\hat{u}_l$) that is optimal to provide, but this does not affect the optimal additional price and the associated distortion. Intuitively, competition acts on the transparent component of the price, but the distortion is generated by the hidden component. The same logic holds for sophisticated-side and naive-side distortions below.

Finally, we note a simple—but for applications important—property of our model with homogeneous distortions: that market outcomes and welfare are unchanged if consumers rather than firms bear the distortionary impact $k(a_l)$, and correctly understand that they do so. A situation in which firms charge f_l, a_l and pay $k(a_l)$ is strategically equivalent to one in which firms charge $f_l - k(a_l)$, a_l and consumers pay $k(a_l)$, so that the two versions of the model feature the same equilibrium additional price, the same utility for naive and sophisticated consumers, and the same profits for firms.

III.C. Sophisticated-Side Distortions

We turn to sophisticated-side distortions, where the distortionary impact falls only on trades with sophisticated consumers. Then:

PROPOSITION 2 (SOPHISTICATED-SIDE DISTORTIONS). Given α, there is a unique symmetric pure-strategy Nash equilibrium, which has additional price $a(\alpha) = (k')^{-1}(\alpha)$. Perfect naïveté-based discrimination ($\alpha_n = 1$, $\alpha_s = 0$) maximizes social welfare. If for $a \in (a(0), a(1))$ the derivative of $\frac{k'(a)}{k''(a)}$ is positive and bounded away from 0, then there is an α^* such that if $\alpha_{ns}, \alpha_n, \alpha_s < \alpha^*$, naïveté-based discrimination strictly lowers welfare.

The condition determining the equilibrium additional price is the same as in the case of homogeneous distortions. Intuitively, all consumers select between products thinking that they will bear the distortion, so they choose as if the distortion was homogeneous. Nevertheless, while perfect naïveté-based discrimination often minimizes welfare for homogeneous distortions, it always maximizes welfare for sophisticated-side distortions. If a firm knows that a consumer is sophisticated—and hence she anticipates any additional price and dislikes the associated cost—it imposes no additional price, so no distortion arises. If a firm knows that a consumer is naive, it can exploit the consumer without triggering a distortion.

Although perfect naïveté-based discrimination has a qualitatively different welfare effect under sophisticated-side distortions than under homogeneous distortions, there is a range in which imperfect naïveté-based discrimination has a similar welfare effect. Namely, if a version of decreasing absolute convexity holds, then a sufficiently small amount of information lowers welfare if the share of naive consumers is sufficiently small. In this case, most consumers in both pools are sophisticated and hence bear a distortionary impact, so welfare is close to that with a homogeneous distortion.

III.D. Naive-Side Distortions

Finally, we consider naive-side distortions, where the distortionary impact arises only on the naive side of the market. The effect of naïveté-based discrimination is particularly simple in this case:

PROPOSITION 3 (NAIVE-SIDE DISTORTIONS). If consumers bear the distortionary impact, there is a unique symmetric pure-strategy Nash equilibrium, which has $a(\alpha) = a_{max}$ for any $\alpha \in (0, 1]$. If firms bear the distortionary impact, there is a unique symmetric pure-strategy Nash equilibrium, which has $a(\alpha) = (k')^{-1}(1)$ for any $\alpha \in (0, 1]$. Both when consumers and when firms bear the distortionary impact, naïveté-based discrimination does not affect social welfare.

If naive consumers bear the distortionary impact, the distortion affects neither a firm's margin nor (since a naive consumer does not anticipate it and a sophisticated consumer does not bear it) a consumer's willingness to accept an offer. As a result, for any

α the firm charges the highest additional price it can ($a = a_{max}$), and naïveté-based discrimination has no effect on welfare. If the firm bears the distortionary impact, then both the benefit and the cost of raising the additional price arises only for naive consumers. As a result, the optimal additional price is again independent of α, and naïveté-based discrimination has no effect on welfare.

III.E. Some Implications for Individual Welfare

Although our main interest is in the implications of naïveté-based discrimination for total welfare, we discuss some implications for individual welfare that are distinct from those of preference-based discrimination. In contrast to our results on social welfare above, the distributional implications depend on the level of competition determined by the product differentiation parameter t. Starting from the symmetric equilibrium, if a firm increases f_l by a small amount, it loses customers either to its rival or to the outside option. We refer to the market as imperfectly competitive in the former case and as monopolistic in the latter case. Whether the industry is imperfectly competitive or monopolistic can depend on α, but for simplicity we assume that naïveté-based discrimination does not affect which case applies.

First, our analysis of the three cases above yields a simple general conclusion:

COROLLARY 1. Under both monopoly and imperfect competition, naïveté-based discrimination is never Pareto-improving. In particular,

i. Under monopoly, naïveté-based discrimination either has no effect on outcomes, or it strictly lowers the welfare of naive consumers in the more naive pool.
ii. Under imperfect competition, naïveté-based discrimination strictly lowers the welfare of sophisticated consumers in the more sophisticated pool.

In a monopolistic market, a sophisticated consumer receives her outside option, so the welfare of a naive consumer is strictly decreasing in the additional price she pays. Since for naive-side distortions naïveté-based discrimination does not affect the additional price, it does not affect any market participant's welfare. But since in the other cases naïveté-based discrimination raises the additional price in the more naive pool, it lowers the welfare of naive consumers in that pool.

Under imperfect competition, an increase in the share of naive consumers—and hence an increase in the profits to be had from the additional price—leads firms to compete more aggressively on the anticipated price. This implies that—similarly to the logic of Gabaix and Laibson (2006) and the literature following it—naive consumers in effect cross-subsidize sophisticated ones, so that a sophisticated consumer's welfare is increasing in the share of naive consumers in her pool. As a result, naïveté-based discrimination lowers the welfare of sophisticated consumers who end up in a pool with fewer naive consumers.

For our second result, we identify circumstances under which naïveté-based discrimination is Pareto-damaging—a possibility that never obtains in basic models of preference-based discrimination. We consider perfect naïveté-based discrimination under homogeneous distortions and imperfect competition. Our proofs imply that due to the above incentive to compete for profitable consumers, an increase in α—much like a common cost reduction does in other Hotelling-type settings—leaves firms' equilibrium profits unchanged. Naïveté-based discrimination therefore also leaves firms' profits unchanged. In addition, Corollary 1 implies that perfect naïveté-based discrimination makes all sophisticated consumers worse off. To complete the picture, we characterize how naïveté-based discrimination affects naive consumers:

PROPOSITION 4. Suppose that the market features a homogeneous distortion, and $v > t + c$.

i. If $k''(a(\alpha))a(\alpha) < 1$ for all α, then for any α_{ns}, perfect naïveté-based discrimination strictly lowers the welfare of naive consumers.
ii. If there is an $\underline{\alpha}$ such that $k''(a(\alpha))a(\alpha) > 1$ for all $\alpha > \underline{\alpha}$, then for any $\alpha_{ns} > \underline{\alpha}$, perfect naïveté-based discrimination strictly raises the welfare of naive consumers.

On the one hand, perfect discrimination eliminates the cross-subsidy from naive to sophisticated consumers, benefiting naive consumers. On the other hand, perfect discrimination leads firms to increase the additional price for naive consumers, hurting naive consumers. The net effect is in general ambiguous, and the latter effect dominates if and only if the proportional responsiveness of the additional price to α is sufficiently high. This responsiveness is $\frac{a'(\alpha)}{a(\alpha)} = \frac{1}{k''(a(\alpha))a(\alpha)}$, giving rise to the conditions in the proposition.

IV. Applications

In this section, we use the insights from our reduced-form framework to study naïveté-based discrimination in models of many markets that researchers have invoked as featuring naive consumers. In all applications, we assume that there are two firms. We specify only consumers' gross utilities from products, but continue to suppose that transportation costs modify these utilities as well as those from the outside option in the same way as above. We say that a market model "simplifies to" our reduced-form model if there is a mapping of the market model's primitives to f, a, v, and k such that the model's equilibrium outcomes and payoffs map to those of the reduced-form model. We organize the applications according to the distortionary impact of naïveté.

IV.A. Homogeneous Distortions

1. Credit. Our leading application is a credit market where partially naive present-biased borrowers underestimate their willingness to pay costly interest on a loan, and firms take advantage of this mistake by lending more than is socially optimal.[17] Lenders interact with consumers over three periods. In period 0, firm l makes an offer consisting of a loan amount b_l, interest rate r_l, and discount d_l to consumers. The discount could, for instance, capture airline miles, cash back, or other credit card perks. If a consumer takes firm l's loan, then in period 1 she chooses an amount $q \in [0, b_l]$ to repay in that period, leaving $(b_l - q)(1 + r_l)$ to be repaid in period 2. Firms' fixed cost of serving a consumer is zero, and they acquire funds at zero interest.

Borrowers have time-inconsistent preferences derived from hyperbolic discounting à la Laibson (1997) and O'Donoghue and Rabin (1999). Self 0's utility from taking firm l's contract is $u(b_l) - q - (b_l - q)(1 + r_l) + d_l$, where $u(\cdot)$ is the gross utility from funds.[18]

17. Our model is a variant of that in Heidhues and Kőszegi (2010), which in turn builds on previous work by DellaVigna and Malmendier (2004) and Eliaz and Spiegler (2006). While the logic of the equilibrium contracts signed by consumers is similar to that in these previous papers, we move beyond the literature in asking how naïveté-based discrimination affects welfare.

18. We assume that self 0 does not discount the cost of repayment relative to the utility from consumption because at the time of accepting a credit offer, typically both are things of the future. Our model also assumes that d_l is given out at the repayment stage (entering utility outside $u(\cdot)$) rather than being added to consumption b_l (in which case it would enter utility inside $u(\cdot)$). After characterizing equilibrium, we provide an endogenous reason for this specification in note 21.

Echoing other uses of a quasilinear framework, our assumption that the disutility from repayment is linear approximates a situation where the loan is taken for a specific good with diminishing marginal utility (e.g., a specific durable good), but repayment is made out of a general budget to which a much less curved indirect utility function applies.[19] We suppose that $u(\cdot)$ is three times continuously differentiable, $u(0) = 0$, $u'(0) \geqslant 1$, $u'(b) > 0$, and $u''(b) < 0$ for all $b \geqslant 0$, and $\lim_{b\to\infty} u'(b) = 0$.

In contrast to self 0, self 1 discounts payments in period 2 by a factor $\beta \leqslant 1$, choosing q to minimize $q + \beta(b_l - q)(1 + r_l)$, and setting $q = 0$ when she is indifferent. To ensure an interior solution to a firm's lending problem, we suppose that $\beta > \frac{1}{2}$ for all consumers.[20] In line with much of the literature on time inconsistency, we take the long-run perspective and equate welfare with self 0's utility. Then, the efficient level of borrowing, b^e, satisfies $u'(b^e) = 1$.

Following O'Donoghue and Rabin (2001), we assume that in period 0 a consumer has point beliefs $\hat{\beta}$ about her future β; that is, she believes that self 1 will choose q to minimize $q + \hat{\beta}(b_l - q)(1 + r_l)$. A consumer chooses a contract or the outside option to maximize her perceived utility, given her prediction about her own future behavior. Firms 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 $\beta = \beta_n < \hat{\beta}$.

We first characterize equilibrium in this model:

LEMMA 1. *The contract (b^*, r^*, d^*) firms offer in a symmetric pure-strategy equilibrium satisfies* $u'(b^*) = 1 - \frac{\alpha(1-\beta_n)}{\beta_n}$, $r^* = \frac{1-\beta_n}{\beta_n}$, *and* $d^* = \max\{\alpha r^* b^* - t, -(u(b^*) - b^*)\}$.

In equilibrium, firms choose the interest rate so that consumers expect to repay their loans in period 1, but naive

19. The same assumption also isolates the overlending distortion generated by the exploitation of consumer naïveté. If the cost of repayment is nonlinear, then additional welfare costs arise because welfare depends not only on the total repayment amount but also on how that amount is distributed between periods 1 and 2. We analyze such a model in Online Appendix B.

20. If $\beta < \frac{1}{2}$, then a borrower is willing to pay more than a 100% interest to delay repayment. Our analysis implies that if all consumers are naive, a lender can then give a discount equal to the loan amount and still run a profit, so that borrowing more seems both costless to the consumer and profitable to the lender. This implies that a lender can always increase profits by increasing the loan amount.

consumers just put off repayment to period 2. Furthermore, to increase the interest naive consumers unexpectedly pay, firms overlend ($b^* > b^e$).[21] This characterization implies that our credit-market model simplifies to our reduced-form model. To state the mapping, we let $x = \frac{\beta_n}{1-\beta_n}$.

LEMMA 2. *Defining* $f = -d$, $a = \frac{b}{x}$, $\overline{a} = \frac{b^e}{x}$, $k(a) = (u(b^e) - b^e) - (u(xa) - xa)$ *for* $xa \geqslant b^e$, *and* $v = u(b^e) - b^e$, *our credit-market model simplifies to the reduced-form model with homogeneous distortions.*

In the credit model, the additional price is the unexpected interest naive consumers pay, and the distortionary impact is the welfare loss from inefficiently high lending. Invoking Proposition 1:

PROPOSITION 5. If the consumption-utility function $u(\cdot)$ satisfies prudence ($u'''(b) \geqslant 0$ for $b \geqslant 0$), then the induced distortionary impact $k(\cdot)$ satisfies decreasing absolute convexity on $[\frac{b^e}{x}, \infty)$, so that for any α_{ns}, α_n, α_s, naïveté-based discrimination strictly lowers social welfare.

Prudence is a commonly assumed property of the consumption-utility function that is equivalent to the precautionary savings motive (Leland 1968). Numerous papers, including empirical studies by Parker and Preston (2005) and Ventura and Eisenhauer (2006) and experimental studies by Deck and Schlesinger (2014) and Noussair, Trautmann, and van de

21. It is worth noting some potentially important implications of our analysis for how firms want to pay the discount d. Since $u'(b^*) < 1$, the consumer derives greater utility from—and hence firm l prefers to disburse—d_l at the repayment stage (entering utility outside $u(\cdot)$) rather than as an addition to consumption b_l (in which case it would enter inside $u(\cdot)$). Furthermore, to avoid lowering the unexpected interest it receives, firm l wants to pay d_l in a way that does not lower the borrower's interest-bearing balance due. These insights both provide an endogenous reason for our (exogenous) specification of how firms pay the discount and are consistent with many forms of real-life credit card perks. Most credit card perks are available only some time after purchase and hence cannot be used to augment the purchase itself. (An exception to this regularity is free rental car insurance.) In addition, possibly as an attempt to avoid decreasing the amount due, many credit card companies dole out perks in goods (e.g., airline miles) that consumers may value less than the cash equivalent. Even when the perk is in the form of cash back, issuers use a variety of incentives to encourage borrowers to use it for purchases rather than debt repayment.

Kuilen (2014), find that the vast majority of the population is prudent. This suggests that in the credit domain, naïveté-based discrimination strictly lowers welfare.

To understand Proposition 5 more deeply and to appreciate the role of prudence, we provide an intuition for the proposition that is specific to this setting. Naïveté-based discrimination increases the extent of overlending to the pool with more naive consumers and lowers the extent of overlending to the pool with more sophisticated consumers. For two reasons the former effect outweighs the latter effect. First, because $u(\cdot)$ is concave, an increase in lending to a consumer hurts social welfare more than an equal decrease in lending to a consumer raises social welfare. Second, prudence implies that the total amount of lending—too high to begin with—increases. For a prudent consumer, the risk of which pool she will be allocated to—and therefore how much she will borrow—increases the expected marginal utility of consumption relative to that with average consumption. Because the same risk does not change the marginal cost of funds or the expected marginal profits from interest payments, average lending increases.[22]

2. Add-Ons 1. Our next application captures situations in which firms charge a high price for a service naive consumers use unexpectedly often, and this induces consumers to undertake socially inefficient ex ante steps to avoid the service. The model potentially applies to any product with add-ons, including bank accounts, mobile phones, and hotels. In Section IV.B we develop another model that potentially applies to any product with add-ons, and in Section V we discuss ways to empirically distinguish the models.

We suppose that firms sell a basic product—such as amusement park rides—with cost c and an additional service—such as at-location food or toys—with cost zero. Firm l chooses $\tilde{f}_l$ for the basic service and $\tilde{a}_l \in [0, \tilde{a}_{max}]$ for the additional service, and consumers observe all prices. Each consumer is interested in buying at most one basic product and can only purchase the additional service from the firm from which she bought the basic

22. If the cost of repayment is nonlinear, then risk also affects the expected marginal cost of lending. Because repayment is in that case split over multiple periods, however, one would expect the curvature of the consumption benefit to matter more than the curvature of the repayment cost, so that prudence still tends to imply that information increases lending. We confirm this intuition in Online Appendix B.

product. The consumer needs to buy the additional service with some probability—for example, if her child breaks down and demands so—but she can take ex ante steps—such as buying another toy to bring along or negotiating with the child—to decrease that probability. Formally, a sophisticated consumer needs the additional service with probability $\theta_s - e$, where θ_s is her baseline probability and e is the ex ante avoidance effort she exerts to lower the probability. Similarly, a naive consumer needs the additional service with probability $\theta_n - e$, where $\theta_n > \theta_s$. Both types initially believe that they will need the additional service with baseline probability θ_s. The cost of avoidance effort is $\kappa(e)$. To ensure that our problem is well behaved, we assume that $\kappa(\cdot)$ is four times continuously differentiable, $\kappa(0) = \kappa'(0) = 0$, $\kappa'(e)$, $\kappa''(e) > 0$ for $e > 0$, $\frac{\kappa'(e)}{\kappa''(e)}$ is strictly increasing in e, and $\tilde{a}_{max} \leqslant \kappa'(\theta_s)$. Finally, the value of the product—whether or not the additional service is needed—is v.

Once again, this model has a reduced form consistent with the framework of Section III. To state the mapping, we denote consumers' optimal avoidance effort given $\tilde{a}_l$ by $e^*(\tilde{a}_l)$.

LEMMA 3. *Defining* $f = \tilde{f} + (\theta_s - e^*(\tilde{a}))\tilde{a} + \kappa(e^*(\tilde{a}))$, $a = (\theta_n - \theta_s)\tilde{a}$ *and* $k(a) = \kappa(e^*(\tilde{a}))$, *our first model of add-ons simplifies to the reduced-form model with homogeneous distortions.*

In this case, the part of add-on expenses that naive consumers fail to anticipate is the additional price, and consumers' cost of avoidance effort is the distortionary impact. Unlike for our credit model, we are unaware of existing evidence that could determine whether $k(\cdot)$ satisfies decreasing absolute convexity. As a potentially informative example, we note that if the cost of avoidance has a power function form with any exponent, naïveté-based discrimination lowers welfare:

EXAMPLE 1. Suppose $\kappa(e) = \phi e^{\gamma}$, with $\phi > 0$ and $\gamma > 1$. Then the induced distortionary impact $k(\cdot)$ satisfies decreasing absolute convexity. Hence, for any α_{ns}, α_n, α_s, naïveté-based discrimination strictly lowers social welfare.

3. Other Applications with Homogeneous Distortions. We mention other applications briefly. First, firms might exacerbate naive consumers' unexpected spending by undertaking distortionary ex ante product modifications that affect all consumers and that therefore generate a homogeneous distortion. For instance, a casino may provide alcohol below cost and create an overly glittery environment to encourage naive consumers to gamble more.[23] Second, when some consumers pay add-on fees unexpectedly often and therefore firms set these fees high, all consumers may worry that they will run into the fees, introducing a homogeneous distortion by lowering all consumers' utility. A traveler who flies with Ryanair, for instance, might be stressed thinking about whether she will manage to avoid all fees.

IV.B. Sophisticated-Side Distortions

1. Add-ons 2. In our second model of add-on pricing, inspired by Grubb (2015), both naive and sophisticated consumers expect to take costly steps to avoid an expensive add-on, but naive consumers fail to do so. We develop the model in the context of mobile phone overage charges, but the same mechanism may apply to bank account overdraft fees and hotel add-ons as well.[24]

A consumer values M minutes of phone usage each month, with the minutes having heterogeneous values $v' \in [0, 1]$. The number of minutes with values less than or equal to v' is $G(v')$, where $G(\cdot)$ is twice continuously differentiable and $G'(v') = g(v') > 0$. Each provider offers a package $f, M - a$, where f is the monthly fee and $M - a \in [0, M]$ is the number of included free minutes (so that a is the number of useful minutes that are excluded). The price of extra minutes is fixed at 1.[25] Independently of how many minutes a consumer calls, a provider's cost of serving her is c. Both naive and sophisticated consumers expect to pay attention and

23. Our working paper, Heidhues and Kőszegi (2014), formalizes this application.

24. The idea that sophisticated consumers pay avoidance costs is discussed by Gabaix and Laibson (2006) in the context of hotels and by Armstrong and Vickers (2012) in the context of bank accounts. One important difference is that—consistent with our arguments in Section III.A for studying naïveté-based discrimination in pools of consumers with the same beliefs—in our setting both naive and sophisticated consumers expect to undertake costly avoidance, whereas in previous work only sophisticated consumers do.

25. Our assumption that the price for extra minutes is exogenously fixed—and sufficiently high for consumers to want to avoid extra minutes—could derive, for instance, from real or implicit regulations on overage fees. Another possible microfoundation for this assumption is that naive consumers' ex post valuation for minutes—which they fail to predict ex ante—is 1. Yet another possibility is that naive consumers follow a heuristic as to whether to call when their free minutes are exhausted, and this heuristic generates a maximum price the provider can charge.

use their phones optimally—that is, only for the $M - a$ minutes of greatest need. Sophisticated consumers indeed pay attention, but naive consumers do not, calling M minutes altogether. Attention itself is costless.[26]

We can once again simplify this model:

LEMMA 4. *Defining* $k(a) = \int_0^{G^{-1}(a)} v'g(v')dv'$ *and* $v = \int_0^1 v'g(v')dv'$*, our second model of add-ons simplifies to the reduced-form model with a sophisticated-side distortion.*

Intuitively, since naive consumers do not avoid overage charges, no distortion arises on the naive side of the market. Sophisticated consumers, on the other hand, do not use the phone for their least valuable a minutes, creating a welfare loss of $\int_0^{G^{-1}(a)} v'g(v')dv'$. Using Proposition 2, perfect naïveté-based discrimination maximizes social welfare.[27]

2. Other Applications with Sophisticated-Side Distortions. Some other applications naturally fit in our category of markets with sophisticated-side distortions. First, firms may offer a rebate that consumers must exert costly effort to cash in, and while all consumers expect to do so, naive consumers forget.[28] Then the savings naive consumers unexpectedly forgo is the additional price, and the effort cost sophisticated consumers pay to return the rebate creates a sophisticated-side distortion. Second, in a dynamic setting firms may increase prices on consumers who automatically renew their contracts, and while all consumers expect to search for better deals to avoid the trap, only sophisticated consumers do.[29]

26. If attention was also costly, sophisticated but not naive consumers would be paying attention costs, generating an additional sophisticated-side distortion.

27. In addition, if $k(\cdot)$ satisfies a version of decreasing absolute convexity, then if α_{ns} is small and the amount of information is small, naïveté-based discrimination lowers welfare. Simple arithmetic shows that $\frac{k'(a)}{k''(a)} = v'g(v')$, where $v' = G^{-1}(a)$, so that decreasing absolute convexity means that $v'g(v')$ is increasing in v'. We are unaware of empirical evidence that can help in determining whether this condition holds.

28. This hypothesis is consistent with experimental evidence on overconfidence about memory by Ericson (2011).

29. Kiss (2014) documents that consumers forgo substantial savings by failing to switch providers in the Hungarian market for mandatory auto liability insurance. Using a structural model, he also estimates that the primary reason for low switching rates is that consumers do not pay attention to the possibility of switching.

Again, the savings naive consumers unexpectedly forgo is the additional price, and the switching cost sophisticated consumers pay is a sophisticated-side distortion.

IV.C. Naive-Side Distortions

1. Useless Extras. We build a model in which firms selling a product can add useless extras that sophisticated but not naive consumers know to avoid. Suppose firms sell a basic product—such as a new car—that costs c to produce and that consumers value at v. A firm charges a base price f for the product, and—intending to buy a basic model—consumers select the firm to buy from based on this base price. Once a consumer arrives, the firm can offer optional extras—such as rust proofing, paint and fabric protection, and prepaid maintenance—that generate no value for consumers.[30] Sophisticated consumers understand that the extras are useless, but naive consumers do not. A package of extras naive consumers think is worth a costs a firm $k(a)$ to produce.

This model again simplifies:

LEMMA 5. *The model of useless extras simplifies to the reduced-form model with a naive-side distortion paid by the firm.*

Here the payment naive consumers make for the extras is the additional price, and the production cost firms pay for the extras is the distortionary impact. By Proposition 3, naïveté-based discrimination does not affect social welfare.

2. Other Applications with Naive-Side Distortions. Again, we briefly discuss other applications. First, consider an insurance market in which a naive consumer is unable to read all the details of her health insurance contract and assumes that she is getting full coverage. As a result, firms add hidden copayments for some contingencies, leading sophisticated consumers—who understand all terms—to buy additional insurance to extend their coverage to full. Because naive consumers forgo valuable insurance but sophisticated consumers do not, this market features a naive-side distortion that is borne by consumers.

Second, in some settings a firm may have to pay customer service, legal, default, or other costs to collect fees that naive consumers unexpectedly owe but are unable or unwilling to pay.

30. The implications are unchanged if consumers have some value for extras, but this value is less than the production cost. The examples above of car dealer extras are widely considered useless; see for instance http://www.forbes.com/2010/12/08/car-buying-warranty-business-autos-dealer-extras_slide.html.

These collection costs generate a naive-side distortion that is borne by the firm.

Third, products with add-ons can generate a naive-side distortion if—instead of underestimating their demand as in Section IV.A or overestimating their avoidance effort as in Section IV.B—naive consumers are unaware of add-on prices. As an example, suppose that firms sell a basic "all-inclusive" holiday package, and at the resort offer extras—such as jet skiing or excursions—that are not included in the basic package. Naive consumers erroneously believe that extras are included, while sophisticated consumers are not interested in extras. Naive consumers' price unawareness brings about an additional price for the resort and can generate a naive-side distortion borne by consumers for multiple reasons. Naive consumers may simply overspend and hence fail to optimally allocate consumption over time, or—if they realize the price of extras at the resort and have a downward-sloping demand curve—their consumption of extras may be distorted.

V. Empirical Issues

In this section, we briefly discuss potential methods for empirically verifying properties of a market that are central in determining the welfare effects of naïveté-based discrimination.

V.A. Verifying Decreasing Absolute Convexity

Our analysis reveals that in a market with a homogeneous distortion, the impact of naïveté-based discrimination hinges on whether $k(\cdot)$ satisfies decreasing absolute convexity (i.e., whether $\frac{k'(a)}{k''(a)}$ is increasing in a). In the spirit of the sufficient-statistics approach to welfare analysis (Chetty 2009), decreasing absolute convexity is (at least in principle) verifiable based on observable market outcomes. Consider a proportional tax of τ imposed either on the additional price or on the total price. A simple derivation shows that in either case $\frac{k'(a(\alpha))}{k''(a(\alpha))} = -(1-\tau)\frac{\partial a(\alpha)}{\partial \tau}$.[31] Hence, if we can observe the responsiveness of the additional price to tax changes at various tax levels, or to a small tax at various levels of α, we know whether $k(\cdot)$ satisfies decreasing absolute convexity. Similar formulas can be obtained for other shocks (e.g., cost shocks) to the industry.

31. In equilibrium, $k'(a(\alpha)) = \alpha(1-\tau)$. Differentiating with respect to τ gives $k''(a(\alpha))\frac{\partial a(\alpha)}{\partial \tau} = -\alpha = \frac{-k'(a(\alpha))}{1-\tau}$. Solving for $\frac{\partial a(\alpha)}{\partial \tau}$ yields the formula.

Unfortunately, we are unaware of existing estimates of the responsiveness of the additional price in any market. But as we note in Section VI.A, a number of researchers have found ways to show that some consumers pay (what we call) additional prices, so there is no reason to believe that the estimates are impossible to obtain. In addition, in a specific application the decreasing absolute convexity of $k(\cdot)$ may be related to an existing property of classical primitives. By Proposition 5, for instance, in the credit model it is implied by a weak condition on the utility function, prudence.

Although it is ultimately an empirical question, some of our analysis suggests that even beyond the credit model, decreasing absolute convexity is a weak condition. For the condition to be violated, $a(\alpha)$ would have to be more responsive to tax changes for higher levels of taxes and lower levels of α—where its starting level is already lower—so that its tax elasticity would have to increase drastically for higher levels of taxes and lower levels of α. This seems to us implausible.[32]

V.B. *Distinguishing Different Models of Add-ons*

As our applications illustrate, the distortion generated in a market with add-ons depends on precisely what mistake naive consumers are making. In the model of Section IV.B and consistent with existing approaches in the literature, naive consumers correctly predict their add-on demand gross of avoidance, but mispredict their avoidance behavior, generating a sophisticated-side distortion. In the model of Section IV.A, naive consumers correctly predict their avoidance behavior but mispredict their add-on demand gross of avoidance, generating a homogeneous distortion. To distinguish these two models empirically, we can test whether consumers correctly predict their avoidance behavior. For instance, how consumers react to a precontractual change in the add-on fee—such as what change in the basic fee makes them indifferent—reveals their expectations about how much of the add-on they will use, and therefore also their expectations about how their avoidance behavior will respond to the fee change.

32. It is, however, easy to give examples of functions that satisfy the properties we have imposed on $k(\cdot)$ in Section III.A but do not satisfy decreasing absolute convexity. Suppose, for instance, that $k''(a)$ increases monotonically on the interval $[a', a'']$, with $k''(a') = 1$ and $k''(a'') = 2$. If $a'' - a'$ is small, then $k'(a)$ changes little over the same interval, so $\frac{k'(a')}{k''(a')} > \frac{k'(a'')}{k''(a'')}$.

Comparing this expectation to their actual response to a surprise postcontractual change in the add-on fee reveals whether the model of Section IV.A or Section IV.B is more appropriate.

In addition, a market with add-ons can generate a naive-side distortion if naive consumers are unaware of add-on prices. This is distinguishable from our other models of add-ons by means of evidence on whether consumers underestimate and are unresponsive to the add-on price. Some evidence is available for bank account overdraft fees, but it is mixed and calls for more research. Consistent with price unawareness but inconsistent with our other models, UK banks believe that demand for bank accounts is unresponsive to overdraft fees (Office of Fair Trading, 2008, paragraph 3.74). Other findings, however, suggest that bank account holders are aware of overdraft fees, although it is unclear whether they understood the fees at the time of account opening. The majority of UK account holders who incurred overdraft charges had heard about such charges beforehand (Office of Fair Trading, 2008, paragraph 4.74), and Stango and Zinman (2014) report evidence that many U.S. consumers try (but often fail) to avoid overdrafting.

VI. Related Literature

VI.A. *Empirical Background*

One of the key assumptions of our model is that naive consumers incur unexpected charges. This assumption is made in different forms in many papers in behavioral industrial organization and is consistent with empirical facts from a number of industries.[33] The other central assumption is that firms acquire and use information about consumer naïveté for designing

33. For instance, Stango and Zinman (2009) find that consumers incur many avoidable fees. Grubb and Osborne (2015) estimate that mobile phone consumers are inattentive to past usage and underappreciate the variance of their own demand. The Office of Fair Trading (2008) reports that most consumers who use overdraft protection do so unexpectedly. Evidence by Agarwal et al. (2008) indicates that many credit card consumers seem to forget or not to know about various fees issuers impose. Ausubel (1991) documents that consumers receiving credit card solicitations overrespond to the introductory ("teaser") interest rate relative to the postintroductory rate, suggesting that they end up revolving debt more than they intended or expected. Regulators are worried about the "bill shock" many mobile phone consumers face when they unknowingly run up charges (Federal Communications Commission 2010). Other work includes Shui and Ausubel (2004) for credit cards, Armstrong and Vickers (2012) for bank accounts,

offers. Although not conclusive, some direct evidence is consistent with this assumption. Gurun, Matvos, and Seru (forthcoming) 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.

Beyond the direct evidence, simple economic logic based on the incentives and information of firms dictates that naïveté-based discrimination is or will soon be pervasive. Since in most of the settings we study naive consumers are more profitable than sophisticated consumers with the same ex ante beliefs and preferences, firms have an incentive to obtain outside information about consumers' naïveté. In addition, researchers have documented several simple correlates of the tendency to make financial mistakes,[34] making it likely that firms also have access to some—perhaps different and probably partial—information regarding naïveté. As emphasized by Bar-Gill and Warren (2008, 23–25), this is especially so given recent technological advances in collecting and processing information about individual consumers. As a simple example, the complexity of the words a person uses in an email message may well be correlated with naïveté, and Google allows firms to condition offers on this information.

VI.B. Related Theory

In asking how outside information about consumers affects economic outcomes, our article is related to the literature on first- and third-degree price discrimination, as well as to the literature on privacy. The existing literature overwhelmingly assumes that the consumer type about which firms may acquire information concerns preferences. To our knowledge, no paper has analyzed the welfare effects of naïveté-based discrimination.

Hall (1997) for printers, and Bucks and Pence (2008) and Gerardi, Goette, and Meier (2009) for mortgages.

34. For instance, Agarwal et al. (2009) find an age pattern in the amount of financial mistakes individuals make, Calvet, Campbell, and Sodini (2007) report that consumers with lower levels of education or income make more investing mistakes, and Stango and Zinman (2011) document that it is possible to predict, based on two simple hypothetical questions on the Survey of Consumer Finances, the consumers who buy the most overpriced loans.

In classical settings, perfect discrimination always maximizes welfare given the number of firms in the market (Stole, 2007). The welfare effect of third-degree preference-based price discrimination, however, is in general ambiguous. Building on a large literature, Aguirre, Cowan, and Vickers (2010) analyze monopolistic third-degree price discrimination and establish how the overall welfare effect depends on the interplay between the misallocation effect first introduced by Pigou (1920) and the output effect originally discussed by Robinson (1933).[35] Bergemann, Brooks, and Morris (2015) show that third-degree price discrimination can generate any combination of producer profit and consumer surplus such that producer profit is at least as high as without information, consumer surplus is nonnegative, and total surplus is at most as high as with efficient trade.

The literature on privacy often finds that it is socially beneficial for firms to know more about consumers or employees. Stigler (1980) argues that the protection of personal information leads firms to substitute other, less efficient forms of information acquisition or screening, and Posner (1981) contends that privacy protection creates asymmetric information that impedes the functioning of markets. Varian (1996) reasons that it is in both a consumer's and a firm's best interest to know which product the consumer would like—this lowers search costs for the consumer—although the consumer would not like the firm to know how much she likes the product.[36]

In considering how firms respond to the presence of naive consumers, our article belongs to the growing literature on behavioral industrial organization.[37] Studying the effects of a different type of naïveté-based discrimination, Johnen (2016) shows that private information about consumer naïveté is valuable even in competitive markets in which private information about consumer preferences is worthless. In addition, complementing our analysis of the welfare effects of outside information about consumer naïveté, several existing papers (e.g., Eliaz and Spiegler 2006, 2008; Heidhues and Kőszegi 2010) study "second-degree naïveté-based

35. Stole (2007) highlights that the same basic logic determines the welfare effects in a homogeneous-good Cournot model, while additional effects are relevant in a model of price competition with differentiated products.

36. See also Taylor (2004), Acquisti and Varian (2005), Calzolari and Pavan (2006), Hermalin and Katz (2006), and Hoffmann, Inderst, and Ottaviani (2013).

37. See Spiegler (2011) for an introduction to and overview of this literature.

discrimination," asking how firms may screen consumers according to naïveté.

VII. Conclusion

Our analysis leaves open several economically important questions regarding the welfare effects of naïveté-based discrimination. As a case in point, all of our analysis ignores possible distortions arising from participation decisions—that consumers respond to the "wrong" prices when deciding whether to buy. Consider, for instance, the perfect-competition limit of our model ($t \to 0$). Then, since the price consumers perceive is below marginal cost, there may be overparticipation in the market. In fact, in Heidhues and Kőszegi (2015) we argue that the participation distortion can be massive. It is easy to see that naïveté-based discrimination can mitigate or exacerbate overparticipation. As an extreme example, suppose that $v < c$, so that no consumer should be served. It may be the case that without discrimination firms cannot break even and hence do not serve any consumer, but discrimination allows firms to serve the more naive pool, lowering welfare. In contrast, it may also be the case that without discrimination consumers are served, but with discrimination the sophisticated pool is no longer profitable to serve, often increasing welfare.

In addition, how naïveté-based discrimination interacts with other forms of discrimination—including third-degree preference-based discrimination and second-degree naïveté- and preference-based discrimination—is an important topic for future research. As a first step in this direction, Online Appendix A studies the interaction between second- and third-degree naïveté-based discrimination in a version of our credit market model. Furthermore, even within our framework, the different types of distortionary impact—homogeneous, sophisticated-side, and naive-side—may interact, and we do not have a general handle on how these interactions depend on information about naïveté. Again focusing on our credit market model, in Online Appendix B we study an empirically plausible case in which multiple distortions are present.

Finally, our article does not analyze potential policy responses to price discrimination. For instance, a commonly advocated solution to privacy concerns is to require firms to obtain a consumer's consent before using her private information. In ongoing work, we investigate whether this policy helps in our framework and find that it does not: because a naive consumer does not understand that the firm will use information to exploit her, she agrees to giving her information too easily. In addition, a regime with required consent may transfer even more money from naive to sophisticated consumers than a regime without required consent.

DICE, University of Düsseldorf
Central European University

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

References

Acquisti, Alessandro, and Hal R. Varian "Conditioning Prices on Purchase History," *Marketing Science*, 24 (2005), 367–381.

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

——, "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation," *Brookings Papers on Economic Activity*, 40 (2009), 51–117.

Aguirre, Iñaki, Simon Cowan, and John Vickers "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 100 (2010), 1601–1615.

Armstrong, Mark, and John Vickers "Consumer Protection and Contingent Charges," *Journal of Economic Literature*, 50 (2012), 477–493.

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

Bar-Gill, Oren, and Elizabeth Warren "Making Credit Safer," *University of Pennsylvania Law Review*, 157 (2008), 1–101.

Benabou, Roland J., and Jean Tirole "Bonus Culture: Competitive Pay, Screening, and Multitasking," *Journal of Political Economy*, 124 (2016), 305–370.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris "The Limits of Price Discrimination," *American Economic Review*, 105 (2015), 921–57.

Bucks, Brian, and Karen Pence "Do Borrowers Know Their Mortgage Terms?," *Journal of Urban Economics*, 64 (2008), 218–233.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy*, 115 (2007), 707–747.

Calzolari, Giacomo, and Alessandro Pavan "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, 130 (2006), 168–204.

Chetty, Raj "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1 (2009), 451–488.

Deck, Cary, and Harris Schlesinger "Consistency of Higher Order Risk Preferences," *Econometrica*, 82 (2014), 1913–1943.

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

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

——, "Consumer Optimism and Price Discrimination," *Theoretical Economics*, 3 (2008), 459–497.

Ericson, Keith Marzilli "Forgetting We Forget: Overconfidence and Memory," *Journal of the European Economic Association*, 9 (2011), 43–60.

Federal Communications Commission, “Comment Sought on Measures Designed to Assist U.S. Wireless Consumers to Avoid ‘Bill Shock’,” Technical Report 2010, CG Docket No. 09-158, Public Notice May 11, 2010.
Gabaix, Xavier, and David Laibson “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 121 (2006), 505–540.
Galperti, Simone “Commitment, Flexibility, and Optimal Screening of Time Inconsistency,” *Econometrica*, 83 (2015), 1425–1465.
Gerardi, Kris, Lorenz Goette, and Stephan Meier “Financial Literacy and Mortgage Outcomes,” Paper presented at the ASSA meetings, San Francisco, CA, 2009.
Grubb, Michael D. “Consumer Inattention and Bill-Shock Regulation,” *Review of Economic Studies*, 82 (2015), 219–257.
Grubb, Michael D., and Matthew Osborne “Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock,” *American Economic Review*, 105 (2015), 234–271.
Gurun, Umit G., Gregor Matvos, and Amit Seru “Advertising Expensive Mortgages,” *Journal of Finance* (forthcoming).
Hall, Robert E. “The Inkjet Aftermarket: An Economic Analysis,” unpublished manuscript, Stanford University, 1997.
Heidhues, Paul, and Botond Kőszegi “Exploiting Naïveté about Self-Control in the Credit Market,” *American Economic Review*, 100 (2010), 2279–2303.
——, “Using Information about Consumer Naïveté to Price Discriminate,” working paper, 2014.
——, “On the Welfare Costs of Naïveté in the US Credit-Card Market,” *Review of Industrial Organization*, 47 (2015), 341–354.
Hermalin, Benjamin, and Michael Katz “Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy,” *Quantitative Marketing and Economics*, 4 (2006), 209–239.
Herweg, Fabian, and Daniel Müller “Price Discrimination in Input Markets: Quantity Discounts and Private Information,” *Economic Journal*, 124 (2014), 776–804.
Hoffmann, Florian, Roman Inderst, and Marco Ottaviani “Hypertargeting, Limited Attention, and Privacy: Implications for Marketing and Campaigning,” working paper, 2013.
Johnen, Johannes “Dynamic Competition in Markets for Deceptive Products,” working paper, ESMT, 2016.
Kiss, András “Salience Switching,” working paper, University of Amsterdam, 2014.
Laibson, David I. “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 112 (1997), 443–477.
Leland, Hayne E. “Saving and Uncertainty: The Precautionary Demand for Saving,” *Quarterly Journal of Economics*, 82 (1968), 465–473.
Noussair, Charles N., Stefan T. Trautmann, and Gijs van de Kuilen “Higher Order Risk Attitudes, Demographics, and Financial Decisions,” *Review of Economic Studies*, 81 (2014), 325–355.
O’Donoghue, Ted, and Matthew Rabin “Doing It Now or Later,” *American Economic Review*, 89 (1999), 103–124.
——, “Choice and Procrastination,” *Quarterly Journal of Economics*, 116 (2001), 121–160.
Office of Fair Trading, “Personal Current Accounts in the UK,” Technical Report, 2008.
Parker, Jonathan A., and Bruce Preston “Precautionary Saving and Consumption Fluctuations,” *American Economic Review*, 95 (2005), 1119–1143.
Pigou, Arthur C. *The Economics of Welfare* (London: Macmillan, 1920).
Posner, Richard A. “The Economics of Privacy,” *American Economic Review*, 71 (1981), 405–409.
Robinson, Joan *The Economics of Imperfect Competition* (London: Macmillian, 1933).

Ru, Hong, and Antoinette Schoar “Do Credit Card Companies Screen for Behavioral Biases?,” working paper, 2016.
Shui, Haiyan, and Lawrence M. Ausubel “Time Inconsistency in the Credit Card Market,” 14th Annual Utah Winter Finance Conference, May 3, 2004. Available at http://ssrn.com/abstract=586622.
Spiegler, Ran *Bounded Rationality and Industrial Organization* (Oxford: Oxford University Press, 2011).
Stango, Victor, and Jonathan Zinman “What Do Consumers Really Pay on Their Checking and Credit Card Accounts? Explicit, Implicit, and Avoidable Costs,” *American Economic Review Papers and Proceedings*, 99 (2009), 424–429.
——, “Fuzzy Math, Disclosure Regulation, and Market Outcomes: Evidence from Truth-in-Lending Reform,” *Review of Financial Studies*, 24 (2011), 506–534.
——, “Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees,” *Review of Financial Studies*, 27 (2014), 990–1030.
Stigler, George J. “An Introduction to Privacy in Economics and Politics,” *Journal of Legal Studies*, 9 (1980), 623–644.
Stole, Lars A. “Price Discrimination and Competition,” in *Handbook of Industrial Organization*, Vol. 3, Mark Armstrong and Robert Porter eds. Amsterdam: Elsevier, (2007), 2221–2299.
Taylor, Curtis R. “Consumer Privacy and the Market for Customer Information,” *RAND Journal of Economics*, 35 (2004), 631–650.
Varian, Hal R. “Economic Aspects of Personal Privacy,” working paper, UC Berkeley, 1996.
Ventura, Luigi, and Joseph G. Eisenhauer “Prudence and Precautionary Saving,” *Journal of Economics and Finance*, 30 (2006), 155–168.