Exploiting Naivete about Self-Control in the Credit Market*

Paul Heidhues
University of Bonn and CEPR

Botond Kőszegi
University of California, Berkeley

October 23, 2007

Abstract

We develop a model of the credit market where competitive lenders offer long-term contracts specifying a non-linear schedule of possible repayment schedules to borrowers who may differ both in their taste for immediate gratification (β) and their prior beliefs about that taste ($\hat{\beta}$). The equilibrium contract to a borrower with given beliefs $\hat{\beta}$ is discontinuous: it features an advantageous option involving fast repayment, and a "penalty" for deviating and repaying later. These predictions seem to resemble some salient features of many credit-card and mortgage contracts. For fully sophisticated borrowers ($\hat{\beta} = \beta$), the discontinuity acts as a commitment device and yields fast repayment. But all other borrowers ($\hat{\beta} > \beta$), even those with an arbitrarily small amount of naivete, end up delaying repayment despite the penalty, and receive little benefit from commitment. Furthermore, because they believe they will repay early, they underestimate the cost of credit, and hence borrow too much. Non-sophisticated borrowers cross-subsidize sophisticated borrowers, while sophisticated borrowers "discipline" firms by preventing very exploitative contracts. We identify natural conditions under which the above results obtain even if firms observe neither β nor $\hat{\beta}$—because all non-sophisticated consumers endogenously choose contracts for which they mispredict their future behavior. Because linear contracts prevent borrowers from discretely misestimating their behavior, requiring credit contracts to have a linear structure can raise social welfare.

*__Preliminary Draft—Exclusively for Circulation at Yale IO Seminar on October 25, 2007. Not for Web Posting!__ We apologize for the underpolished nature of many parts of this paper, and for missing proofs.

1 Introduction

A growing empirical literature argues that a time-inconsistent taste for immediate gratification, possibly in combination with naivete about this taste, is necessary to explain consumer behavior with respect to savings and credit.[1] Probably motivated by similar intuitions regarding borrowers' behavior, numerous policymakers and writers in the popular press, as well as academics, have raised concerns that some debt contracts may be harmful to consumers. Durkin (2000) reports that 51% of consumers consider using credit cards as "bad," and 82-93% think credit-card companies "make too much credit available," "mislead a lot of people," or "make it hard for people to get out of debt." Similarly, Engel and McCoy (2002) argue that some features of subprime mortgages are designed to exploit naive consumers. Yet these empirical and "consumer-protection" literatures remain largely disconnected: the former one does not consider the welfare effects of credit contracts, while the latter one is not built on a solid theoretical foundation.

To bring rigorous economic analysis to bear on borrowing and consumer well-being, we adapt theories by DellaVigna and Malmendier (2004), Kőszegi (2005), and Eliaz and Spiegler (2006) to credit markets, extend these theories by allowing heterogeneity in the taste for immediate gratification as well as naivete, and analyze the welfare consequences of credit contracts in detail. We replicate a key prediction of all three papers, that the price of delaying repayment beyond the baseline—e.g. by revolving debt on a credit card or by refinancing a mortgage—will be above the lender's marginal cost. But our model also explains another pervasive feature of many credit contracts, a discontinuous penalty for failing to follow the favorable repayment schedule. For instance, most credit cards do not charge interest on *any* purchases if the borrower repays the full balance each month, but—although the lender's cost is presumably continuous—do charge interest on *all* purchases if she revolves even $1. Similarly, subprime mortgages typically feature large refinancing penalties. We show that these discontinuities in turn have important welfare consequences. *Per-*

[1] See, for instance, Ausubel (1991), Laibson, Repetto, and Tobacman (1998), Ausubel (1999), Angeletos, Laibson, Repetto, Tobacman, and Weinberg (2001), Bertaut and Haliassos (2002), DellaVigna and Malmendier (2004), Shui and Ausubel (2004), and Laibson, Repetto, and Tobacman (2007). Bolstering this literature is empirical research by DellaVigna and Paserman (2005) and Paserman (2006) on job search, Fang and Silverman (2007) on welfare-program participation, DellaVigna and Malmendier (2006) on health-club attendance, and Oster and Scott-Morton (2005) on magazine subscription prices, which finds evidence of a taste for immediate gratification in these other domains.

fectly sophisticated borrowers—those who understand their time inconsistency perfectly—benefit because the discontinuous penalty acts as a commitment device that keeps repayment on schedule. In contrast, all other borrowers, even those with an *arbitrarily small* amount of naivete regarding their time inconsistency, choose a contract in which they delay repayment despite the penalty, receiving little benefit from commitment. Worse, because they falsely believe the penalty provides sufficient commitment to guarantee timely repayment, they tend to borrow too much, lowering their welfare further. Hence, a policy of disallowing discontinuous penalties, for example by requiring credit contracts to have a simple linear structure, would raise social welfare.

Section 2 presents our model, which is a modification of Kőszegi (2005) and Eliaz and Spiegler (2006) in a number of ways.[2] There are three periods, 0, 1, and 2. In period 0, the person has total utility $c - k(q) - k(r)$, where c is the amount she borrows and consumes, q and r are the amounts she repays in periods 1 and 2, and $k(\cdot)$ is a strictly increasing and strictly convex function representing the cost of repayment. Self 1, the period-1 incarnation of the individual, acts to maximize $-k(q) - \beta k(r)$ for some $0 < \beta \leq 1$. Hence, for $\beta < 1$, the consumer has a time-inconsistent taste for immediate gratification: the period-1 cost of repayment weighs more heavily in self 1's decision than the consumer would have preferred earlier. Consistent with much of the literature, we take the long-term perspective and equate the consumer's welfare with self 0's utility. To allow for self 0 to be overoptimistic regarding her future taste for immediate gratification, we assume that she believes she will maximize $-k(q) - \hat{\beta} k(r)$ in the future, so that $\hat{\beta}$ satisfying $\beta \leq \hat{\beta} \leq 1$ represents her beliefs about β.

The consumers introduced above contract with competitive suppliers of credit, who have funds available to them at an interest rate of zero. In the bulk of the paper, we consider situations where firms and consumers can sign arbitrary non-linear long-term contracts in period 0. A credit contract specifies a consumption level c as well as a menu of installment plans (q, r) from which the

[2] Eliaz and Spiegler develop a two-period model in which a monopolist offers contracts to a population of consumers with homogeneous ex-ante and ex-post preferences about an action to be taken in the second period, but heterogeneous beliefs about the change in preferences. Like Eliaz and Spiegler, we take a setting where firms can offer contracts to consumers whose tastes change in a way they may not anticipate. But by considering heterogeneity in preferences in addition to beliefs, by assuming a different form of naivete about preferences, and by focusing on perfect competition, many of our general results are different from theirs. And by considering a more specific setting developed for the credit market, we are able to derive more features of the equilibrium contracts.

borrower can choose in period 1. We define a competitive equilibrium as a set of contracts such that each contract earns zero profits, and no firm can offer a contract that some consumers accept and that earns strictly positive profits. To simplify our statements and exposition, we also require that each feature of each contract be relevant for either the firm's or some type of consumer's evaluation of the contract.

Both as illuminating points of comparison and as possible policy interventions, we also consider markets where contracts are restricted to have a linear structure, with an interest rate determining the cost of delaying repayment in period 1. Since non-linear contracts provide an opportunity to commit fully to a repayment schedule ex ante but linear contracts do not, it would seem that this restriction should lower welfare. In a *long-term restricted* market, firms and consumers can still sign contracts in period 0 that determine the future interest rate for delaying repayment. This arrangement still allows for some commitment through the strategic choice of the interest rate. In a *short-term* market, only one-period contracts are possible. Since this arrangement leaves the period-1 borrowing decision entirely in self 1's hands, it seems the worst of all.

We begin in Section 3 by assuming that β and $\hat{\beta}$ are known to firms, considering first the case of $\hat{\beta} = \beta$, perfect sophistication. Since the borrower correctly predicts her own behavior ex ante, the only contract that survives competitive pressure is the one that maximizes her ex-ante utility: it specifies the optimal consumption level, and commits her to repay in equal installments. A sophisticated consumer therefore benefits from the commitment made possible by long-term contracting.

The fully committed repayment schedule, however, is not a competitive equilibrium when consumers are not perfectly sophisticated—when $\hat{\beta} > \beta$. If this was the only contract in the market, a creditor could offer a slightly higher consumption level without increasing either payment, but also include an option for the consumer to delay part of the repayment for a non-trivial fee. Since the customer's period-1 demand for credit is at least slightly higher than foreseen, a firm can design the penalty to induce unexpected switching, so the contract is both attractive to the borrower and profitable to the firm. Because in equilibrium the borrower changes her mind regarding repayment—so that only self 1 cares about the exact terms of the chosen option—just as in the short-term market the installment plan she chooses in the end caters entirely to self 1's preferences. Hence, the ability to commit to a repayment schedule does not help in overcoming the consumer's taste for immediate gratification at all. Even worse, because the consumer mispredicts her behavior, she underestimates the cost of credit, and borrows too much. Hence, for any non-sophisticated consumer, *no matter how close to sophisticated*, the ability to write non-linear contracts lowers welfare even relative to the short-term market.

Given the above problem, a restricted long-term market often Pareto-dominates the unrestricted market. By choosing a contract with gross interest rate $1/\beta$, sophisticated borrowers can still repay in equal installments and achieve the highest possible long-term utility. More importantly, if a non-sophisticated borrower is not too naive, a linear contract prevents her from seriously mispredicting her future behavior, and hence raises her utility. Even if many consumers are very naive, a restricted market with an interest-rate cap, which in effect limits the misprediction of the creditors' repayment costs, is better than the unrestricted market.

To demonstrate the robustness of the above insights, as well as to elaborate on some important issues, in Section 4 we consider equilibria where firms may not know β or $\hat{\beta}$. We begin by assuming that $\hat{\beta}$ is observed but β is not, and in the text assume that there are two types of borrowers for a given $\hat{\beta}$, one sophisticated ($\beta = \hat{\beta}$) and one non-sophisticated ($\beta < \hat{\beta}$).[3] Since a borrower's behavior in period 0 depends only on $\hat{\beta}$, all consumers with a given $\hat{\beta}$ accept the same contract. This contract features a repayment schedule with a higher installment in period 1 than in period 2, and a costlier option catered to self 1's taste for immediate gratification. Confirming the discontinuity above, sophisticated borrowers choose the former option, but non-sophisticated ones, even if close to sophisticated, the latter one. As a result, non-sophisticated borrowers have lower welfare than in the short-term market. The inefficiently front-loaded repayment schedule serves to increase the profits a firm can extract from non-sophisticated consumers' unexpected willingness to pay to delay repayment. We argue that these features of equilibrium contracts (an inefficiently front-loaded schedule along with expensive options to switch out of it) resemble some real-life features of credit-card and mortgage contracts.

[3] In the appendix, we consider more general distributions.

We also consider how the mixture of sophisticated and non-sophisticated borrowers affects social welfare as well as each type's individual welfare. As in the models of DellaVigna and Malmendier (2004) and Gabaix and Laibson (2006), competitive firms make money on non-sophisticated consumers and lose money on sophisticated consumers. As a result of this cross-subsidy, an increase in the proportion of non-sophisticated consumers always benefits sophisticated consumers. Similarly, since this increase means there are fewer sophisticated consumers to cross-subsidize and more consumers to cross-subsidize them, there is a force for the increase to benefit non-sophisticated consumers as well. In our model, however, there is a completely different, off-setting, force. As the proportion of non-sophisticated consumers increases, firms shift their strategy from serving sophisticated consumers toward exploiting non-sophisticated consumers, decreasing the latter's welfare. In this sense, sophisticated consumers provide market discipline by preventing firms from offering very exploitative contracts.

The above cross-subsidy, however, means that intervention in the market has slightly different implications than when β is known. Because a restricted long-term market can substantially decrease the cross-subsidy, it hurts sophisticated consumers, and hence is not Pareto-improving. But if the consumer is not too naive, the restricted market still increases social welfare because it brings the distorted consumption and repayment terms above much closer to optimal.

Finally, we assume that firms know neither β nor $\hat{\beta}$, and impose (roughly) that if a borrower is more optimistic about her future tastes (she has a higher $\hat{\beta}$), she tends to be less sophisticated. We show that in the unique competitive equilibrium of this market with two-dimensional heterogeneity, the above contracts separate consumers first according to $\hat{\beta}$ in period 0, and then according to β in period 1. Since firms make money on non-sophisticated consumers, they compete more fiercely for consumers with a higher $\hat{\beta}$, so that the terms of the credit contract when following the ex-ante preferred early repayment schedule are improving in $\hat{\beta}$. Consumers, looking for the most favorable contract for which they believe they will follow the preferred repayment schedule, therefore choose a contract corresponding to their $\hat{\beta}$. As a result, all non-sophisticated consumers endogenously select a contract under which they *do* delay the bulk of repayment to period 2, so that their behavior and welfare is discontinuously different from that of sophisticated consumers.

This setting allows us to ask whether a borrower benefits from becoming more sophisticated. By simple revealed-preference considerations, it is clear that the borrower is best off if she is fully sophisticated ($\hat{\beta} = \beta$). If she becomes more sophisticated but falls short of becoming fully sophisticated, however, her welfare will often *decrease*. As she becomes more pessimistic about her future taste for immediate gratification, she chooses a credit contract with which it is easier to stick to the original repayment schedule. But since this contract is worse of a deal up front and she ends up changing her mind about repayment anyhow, she is worse off.

2 A Model of the Credit Market

2.1 Setup

In this section, we introduce our model of a credit market where borrowers have time-inconsistent preferences regarding the repayment schedule and may have an incomplete understanding of these preferences. There are three periods, $t = 0, 1, 2$, with consumption and borrowing set in period 0 and repayment occurring in periods 1 and 2. Consuming an amount $c \geq 0$ generates utility c, and repaying amounts $q \geq 0$ and $r \geq 0$ in periods 1 and 2 implies instantaneous utility costs of $k(q)$ and $k(r)$, respectively. Although we will throughout refer to q as a repayment choice, decreases in q can also be interpreted as extra borrowing rather than less repayment, such as when a person takes a home equity loan on her mortgaged house for further consumption. $k(\cdot)$ is a twice-differentiable cost function with $k(0) = 0, k'(0) \geq 0$, and $k''(x) > 0$ for all $x \geq 0$. We assume that $k'(0)$ is sufficiently low for first-order conditions below to describe optimal contracts.

The crucial assumptions of our model are on how the creditor makes intertemporal decisions. Our formulation is motivated by a time-inconsistent taste for immediate gratification as modeled in Strotz (1956) and Laibson (1997), but is also consistent with a cue-based overweighting of present consumption in the spirit of Bernheim and Rangel (2004). Self 0's utility is $c - k(q) - k(r)$, and self 1 instead maximizes $-k(q) - \beta k(r)$, where $0 < \beta \leq 1$, and β parameterizes the taste for immediate gratification.[4] While realistically a taste for immediate gratification applies to all periods, for several

[4] Since self 1 makes no decision regarding consumption, our analysis would be unaffected if her utility also included a consumption term.

reasons we assume that self 0 does not discount the cost of repayment by a factor of β. Many types of borrowing motivating our analysis—such as first mortgages and credit-card purchases of durable goods—are (largely) for future consumption, and in this case self 0 is likely to weight consumption and repayment close to equally.[5] Our results highlight that excessive borrowing results even in this case. Of course, if self 0 discounted repayment costs relative to the utility from consumption, the overborrowing would be exacerbated. Even in that case, our specification is useful to isolate how future preferences and beliefs about those preferences affect borrowing and repayment holding current preferences constant.

In the spirit of much of the literature on time inconsistency (DellaVigna and Malmendier 2004, Gruber and Kőszegi 2004, O'Donoghue and Rabin 2006, for example), we equate welfare with long-run, period-0 preferences, and put zero weight on self 1's taste for immediate gratification. Although we simplify things by considering a three-period model, in reality time inconsistency seems to be mostly about very immediate gratification that plays out over many short periods. Hence, arguments by O'Donoghue and Rabin (2006) in favor of a long-run perspective apply: for any particular month of a person's life, only the self alive in that month discounts future months heavily, so it seems reasonable to put most welfare weight on the many earlier selves who instead prefer a more equal weighting of that month and later months. In addition, the model in Bernheim and Rangel (2004) can be interpreted as saying that a taste for immediate gratification is often a mistake not reflecting true welfare.

The main purpose of this paper is to investigate borrowing and repayment outcomes as a function of the borrower's sophistication regarding her future taste for immediate gratification. A *sophisticated* self 0 fully realizes that she will have different preferences in the future, and optimizes taking into account self 1's correctly anticipated behavior. At the other extreme, a *naive* self 0 is completely unaware that she will be more impatient in the future than she would now like. She simply makes plans to maximize her utility $c-k(q)-k(r)$ and takes the period-0 step in the optimal plan, not realizing that in period 1 she may prefer to change her mind. To allow us to address

[5] Indeed, although high credit-card borrowing is often invoked as indicating a time-inconsistent taste for immediate consumption, existing evidence indicates that most credit-card spending is on durables and other future-oriented goods.

how much sophistication or naivete is required to generate certain outcomes, we also consider intermediate levels of sophistication. Following O'Donoghue and Rabin (2001), we suppose that self 0 believes with certainty that self 1's utility function will be $-k(q) - \hat{\beta}k(r)$.[6] The parameter $\hat{\beta}$ reflects self 0's beliefs about β, so that $\hat{\beta} = \beta$ corresponds to perfect sophistication, whereas $\hat{\beta} = 1$ corresponds to complete naivete, and more generally $\hat{\beta}$ is a measure of sophistication. In period 0, the borrower acts to maximize her utility given these beliefs about her future preferences. We assume that there are finitely many possible β's and $\hat{\beta}$'s in the population.

Our formalization of partial naivete is substantially different from that of Eliaz and Spiegler (2006) and Asheim (2007). Translated into our model, they assume that a consumer has probabilistic beliefs about her future β, putting some probability q on being time-consistent ($\beta = 1$) and probability $1 - q$ on the true β. In their formulation, it is q that measures the degree of naivete. Using a specification of consumer beliefs as a full distribution over β (which subsumes both the O'Donoghue-Rabin and Eliaz-Spiegler models as special cases), Heidhues and Kőszegi (2007) show that the probability the consumer assigns to the true β or lower is typically crucial in determining her behavior and welfare in a discrete choice problem. If this probability is low, the consumer arranges for too little self-control, and hence later breaks down. Since the optimal contracts in our setting induce a kind of discrete choice problem, what seems to matter for our results is that a non-sophisticated borrower assigns zero (or low) probability to her taste for immediate gratification being as low as it actually is.

The borrowers introduced above can write credit contracts with competitive, risk-neutral, profit-maximizing firms. To simplify exposition, we assume that firms face an interest rate of zero, and that there is no possibility of default.[7] For the bulk of our analysis, we assume that firms and consumers can write exclusive non-linear contracts in period 0 regarding the consumption and repayment schedule, and once a consumer signs a contract with a firm, she cannot interact with other firms. We discuss the role of such full exclusivity in Section 3.2, arguing that the logic of our results holds under much milder conditions.

[6] In particular, this means that she does not make inferences about β from the contracts offered to her in the market.

[7] For models of market equilibria when consumers can default (but have time-consistent preferences and are rational), see Livshits, MacGee, and Tertilt (2006, 2007).

An *unrestricted credit contract* is defined as an initial consumption c along with a finite menu $C = \{(q_s, r_s)\}_{s \in S}$ of repayment options, and is denoted by (c, C). We think of both consumer behavior in period 1 and perceptions about that behavior in period 0 as being determined by an incentive-compatible map:

Definition 1. A map $\beta \mapsto (q(\beta), r(\beta)) \in C$ is *incentive compatible* if for each β,

$$k(q(\beta)) + \beta k(r(\beta)) \leq k(q) + \beta k(r) \text{ for all } (q, r) \in C.$$

If the consumer accepts the contract, she believes in period 0 that she will choose $q(\hat{\beta}), r(\hat{\beta})$, and then she will actually choose $q(\beta), r(\beta)$. To enable us to focus on the contracts accepted by consumers and to simplify our statements regarding the features and uniqueness of equilibrium, we suppress the strategic interaction between firms and define equilibrium directly in terms of the contracts that survive competitive pressure:

Definition 2. A *competitive equilibrium* is a set of contracts $\{(c_i, C_i)\}$ and incentive-compatible maps $q_i(\beta), r_i(\beta)$ defining perceptions and behavior satisfying the following properties:

1. [Non-redundancy] For each (c_i, C_i) there is a type $(\hat{\beta}, \beta)$ who chooses (c_i, C_i) from the offered contracts. For each installment plan $(q_j, r_j) \in C_i$, either $(q_j, r_j) = (q_i(\hat{\beta}), r_i(\hat{\beta}))$, or $(q_j, r_j) = (q_i(\beta), r_i(\beta))$, or there is a type $(\hat{\beta}', \beta')$ who does not choose the contract (c_i, C_i) but would do so if the option (q_j, r_j) was omitted.

2. [Competitive market] Each (c_i, C_i) yields zero expected profits.

3. [No profitable deviation] There is no contract (c', C') and incentive-compatible map $q'(\beta), r'(\beta)$ defining perceptions and behavior such that (i) some types strictly prefer (c', C') over any (c_i, C_i); and (ii) given the types who strictly prefer (c', C'), this contract yields positive expected profits.

Our first requirement for competitive equilibrium is a non-redundancy condition: each contract is accepted by some consumer, and all the repayment options in each contract are relevant in that they affect consumer expectations or behavior. More precisely, a repayment option is non-redundant if some consumers expect to choose it or do choose it, and if eliminating this option would affect some consumers' contract choice. Our other two conditions are typical for competitive situations, saying that firms earn zero profits by offering these contracts, and that firms can do no better.

For our analysis of the above market with non-linear contracts, we will use two benchmarks in both of which contracts are restricted to have a linear structure. These benchmarks serve both as conceptually useful comparisons for some of our welfare results, and, to the extent that linear contracts generate higher social welfare than more complicated contracts, they suggest a simple market intervention. In the first benchmark, which we will call the "long-term restricted" market, firms and consumers can still write exclusive contracts in period 0, but the repayment options that are allowed must be linearly related: $q + Rr$ must be a constant for some R. . Hence, firms can set the gross interest rate R, but given the agreed-upon interest rate, borrowers can allocate repayment as they wish. This imposes a simple restriction on possible contracts while maintaining the basic timing structure, so it seems like a potentially feasible market intervention.[8]

In the second benchmark, which we will refer to as a "short-term" market, competitive lenders offer loans in periods 0 and 1 to be repaid in the next period. Clearly, the interest rate in this market will be zero. Because it is both less feasible and (as will be clear from our analysis) worse than the long-term restricted market, we do not think of the short-term market as a policy option. But it is a theoretically interesting extreme case: whereas the long-term restricted market still allows for some commitment through the choice of R, a competitive short-term market offers no self-control benefits whatsoever. Because the interest rate is competitive and self 1 is completely free to allocate repayment across periods 1 and 2, the repayment schedule is determined entirely by her preferences. In this light, our key welfare results are quite surprising: in the unrestricted long-term market, where full commitment to a consumption level is feasible, non-sophisticated consumers do worse than in the short-term market, where no commitment is possible at all.

Of course, there are many possible benchmarks to which we could compare equilibrium outcomes, and, even more importantly, also many plausible market interventions we could contemplate. For instance, because in our model all consumers know their long-term preferences in period 0, re-

[8] Strictly speaking, we have defined a competitive equilibrium only for the case of unrestricted contracts. When considering the restricted long-term market, one needs to replace set of contracts $\{(c_i, C_i)\}$ by the set of admissible linear contracts.

quiring them to commit fully to a repayment schedule would maximize long-term utility. This intervention, however, is clearly suboptimal if consumers are subject to ex-post shocks in their financial circumstances. More generally, because our model abstracts from many considerations that would be relevant for policy, doing a full-fledged analysis of optimal policy does not seem very useful. More narrowly, our points below about linear contracts demonstrate merely that eliminating discontinuities in contract terms can be welfare-increasing. This point seems robust to additional features of the credit market missing from our model.

2.2 Restating the Problem

As a preliminary step to our analysis, we identify a way to restate the requirements of a competitive equilibrium (Definition 2) in contract-theoretic terms when $\hat{\beta}$ is known. This will allow us to use contract-theoretic methods throughout the paper, and will help clarify the economic forces behind equilibrium contracts.

Consider the economic situation created by our model from the point of view of a single firm: given the other contracts available to consumers in the market, what is the best the firm can do? Let $\underline{u}$ be a consumer's perceived utility from the perspective of period 0 if she accepts the perceived best alternative contract offered in the market. For any contract our firm may offer, there is a repayment schedule the consumer will expect to choose given her period-0 beliefs $\hat{\beta}$, and there is an option she will actually choose given her true β. Hence, we can think of the firm as choosing consumption c along with the former "decoy" repayment option and the latter "chosen" repayment options subject to the following constraints. First, self 0 must be willing to accept the firm's offer over her best alternative: self 0's utility from the decoy option must be at least $\underline{u}$. This is a version of the standard participation constraint (PC), except that self 0 may make her participation decision based on incorrectly forecasted future behavior. Second, if self 0 is to think that she will choose the decoy option, then given her beliefs $\hat{\beta}$ she must think she will prefer it to the other available options. We call these constraints the perceived-choice constraints (PCC). Third, if a consumer with short-term impatience β is to actually choose the repayment schedule intended for her, she has to prefer it to the other repayment options. This is analogous to standard incentive-compatibility constraints (IC) for self 1.

It is clear that a competitive-equilibrium contract must be a solution to the above maximization problem with $\underline{u}$ defined as self 0's perceived expected utility from accepting the contract: if this was not the case, a firm could solve for the optimal contract and increase c slightly, attracting all consumers and making strictly positive expected profits. In addition, for the solution to the above maximization problem to be a competitive equilibrium, $\underline{u}$ must be such that the highest achievable expected profit is zero. In fact, this is also sufficient:[9]

Lemma 1. *Suppose $\hat{\beta}$ is known, and the possible β's are β_1, β_2. The contract with consumption c and repayment options $\{(q(\hat{\beta}), r(\hat{\beta})), (q(\beta_1), r(\beta_1)), (q(\beta_2), r(\beta_2))\}$ is a competitive equilibrium if and only if there is a $\underline{u}$ such that the contract maximizes expected profits subject to a PC with perceived outside option $\underline{u}$, PCC's, and IC's, and the maximum profit level is zero.*

3 Non-Linear Contracting with Known β and $\hat{\beta}$

We now turn to analyzing market outcomes when lenders and borrowers can sign non-linear contracts in period 0, beginning in this section with the case when both β and $\hat{\beta}$ are known to firms. Section 3.1 presents the results, and Section 3.2 discusses how they would change with alternatives to our assumption of fully exclusive contracts. In Section 4, we demonstrate that many of our key conclusions survive when instead β and $\hat{\beta}$ are unknown to firms.

3.1 Analysis

Our point of departure is the simple remark that if consumers are time consistent, they do not prefer a non-linear contract over a linear one:

Fact 1. *If $\hat{\beta} = \beta = 1$, the competitive-equilibrium contract replicates the consumption and repayment outcomes that obtain with the short-term market and the restricted long-term market.*

For the rest of the section, we assume that $\beta < 1$. First, we consider the (easier) case of a perfectly sophisticated consumer, for whom $\hat{\beta} = \beta$. Whatever choice set of repayment options such

[9] While the result is true for any finite number of β's, we state it here for two β's because that is all we will use in the text.

a consumer is offered, she knows the option she will choose in the end. Hence, applying Lemma 1, the decoy option is identical to the chosen option, there is only one repayment option in a non-redundant contract, and the PCC and IC constraints in the firm's optimization problem are trivial. A competitive-equilibrium contract therefore solves

$$\max_{c,q,r} \quad q + r - c$$
$$\text{s.t.} \quad c - k(q) - k(r) \geq \underline{u}. \tag{PC}$$

It is clear that PC is satisfied with equality; otherwise, the firm could increase profits by lowering c. Plugging PC into the maximand, we can rewrite the firm's problem as

$$\max_{q,r} \quad q + r - k(q) - k(r).$$

This leads to the following proposition:

Proposition 1. *Suppose β and $\hat{\beta}$ are known, and $\hat{\beta} = \beta$. Then, the competitive-equilibrium contract has a single repayment option satisfying $k'(q) = k'(r) = 1$, and $c = q + r$. The consumer's utility is strictly higher than in the short-term market.*

To maximize what it can charge a sophisticated consumer, a firm offers self 0—the self who signs the contract—the option that maximizes her utility, including a repayment plan that commits her to repay in two equal installments.

The situation is entirely different for a non-sophisticated borrower, for whom $\hat{\beta} > \beta$. A non-redundant contract consists of a consumption level c, a decoy repayment schedule $(\hat{q}, \hat{r})$ self 0 expects to choose, and a repayment schedule (q, r) self 1 actually chooses—with the optimal solution determining whether $(q,r)=(\hat{q},\hat{r})$. By Lemma 1, these solve

$$\max_{c,q,r,\hat{q},\hat{r}} \quad q+r-c$$
$$\text{s.t.} \quad c-k(\hat{q})-k(\hat{r}) \geq \underline{u}, \tag{PC}$$
$$-k(\hat{q})-\hat{\beta}k(\hat{r}) \geq -k(q)-\hat{\beta}k(r), \tag{PCC}$$
$$-k(q)-\beta k(r) \geq -k(\hat{q})-\beta k(\hat{r}), \tag{IC}$$

As before, PC binds because otherwise the firm could increase profits by reducing c. In addition, IC must bind: if self 1 strictly preferred the chosen option, a firm could deviate by slightly decreasing $\hat{q}$ and slightly increasing q, still satisfying PC and PCC and increasing profits.

Given that IC binds and $\hat{\beta}>\beta$, PCC is equivalent to $q \leq \hat{q}$. Intuitively, if the real self 1 is indifferent between two repayment options, then a more patient self 1—which self 0 believes she will be—would prefer the option with more repayment early. Conjecturing that the firm will want to set q and $\hat{q}$ satisfying $q \leq \hat{q}$, we ignore PCC, and (unsurprisingly) find that it is satisfied in the solution to the relaxed problem.

The relaxed problem is

$$\max_{c,q,r,\hat{q},\hat{r}} \quad q+r-c$$
$$\text{s.t.} \quad c-k(\hat{q})-k(\hat{r}) = \underline{u}, \tag{PC}$$
$$-k(q)-\beta k(r) = -k(\hat{q})-\beta k(\hat{r}). \tag{IC}$$

Notice that in the optimal solution, $k(\hat{r})$ must be zero: otherwise, the firm could decrease $k(\hat{r})$ and increase $k(\hat{q})$ by the same amount, leaving PC unaffected and creating slack in IC, allowing it to increase q. Using this, we can express $k(q)$ from IC and plug it into PC to get

$$c=k(q)+\beta k(r)+\underline{u}.$$

Plugging c into the firm's maximand leads to the characterization of the equilibrium contract in

Proposition 2 below. Before stating the result, we introduce an assumption we will use to compare welfare in the unrestricted market to that in the short-term market.[10]

Assumption 1. One of the following two assumptions hold:

I. [k is a power function.] $k(x) = x^\rho$ with $\rho > 1$.

II. [k is derived from a CRRA utility function with coefficient of relative risk aversion at least 1.] For some $y > 0, \rho > 0$, $k(x) = (y - x)^{-\rho} - y^{-\rho}$.

We are now ready to characterize the competitive-equilibrium contract:

Proposition 2. *Suppose β and $\hat{\beta} > \beta$ are known. Then, the competitive-equilibrium contract has two repayment options, with the consumer expecting to choose $\hat{q} > 0, \hat{r} = 0$, and actually choosing q, r satisfying $k'(q) = 1, k'(r) = 1/\beta$. Consumption is $c = q + r$, higher than in the short-term market, and higher than that of a sophisticated consumer. The consumer strictly prefers to sign the equilibrium non-linear contract over participating in the short-term market. But if Assumption 1 holds, she has strictly lower utility in the long-term unrestricted market than in the short-term market.*

The first important feature of the equilibrium contract is that it is flexible in a way that induces the borrower to unexpectedly change her mind regarding repayment. Intuitively, the flexible contract offers an option for the borrower to repay relatively little, but also introduces an option to delay part of the repayment for a premium. Thinking that she will not use the option, the consumer likes the deal. But since she changes her mind, the firm earns higher profits than with a committed contract.

In addition, since $k'(q) = \beta k'(r)$, self 1's preferences fully determine the repayment schedule eventually chosen by the borrower, so that the ability to write long-term contracts, and hence the ability to commit perfectly to a repayment schedule, does not mitigate the consumer's time inconsistency regarding repayment at all. Intuitively, once the firm designs the contract to induce

[10] Assumption 1 ensures that in the short-term market—which generates a classical consumption-savings problem with hyperbolic discounting analyzed by Laibson (1997)—borrowing is increasing in naivete. Our welfare arguments will make use of this feature by comparing outcomes and welfare with non-linear contracts to that of fully naive consumers in the short-term market. Since Assumption 1 allows for a wide class of cost functions, including one derived from the CRRA utility function with coefficient of relative risk aversion greater than 1, it seems unproblematic.

unexpected switching, its goal with the chosen option is to maximize the gains from trade with self 1, so it designs the option with self 1's behavior in mind.

The equilibrium non-linear contract not only matches the short-term market in catering repayment fully to self 1, it also induces more borrowing than in the short-term market, and hence (if Assumption 1 holds) leads to lower welfare than even in that market.[11] In other words, even *given* that repayment will be made according to self 1's preferences, consumption is higher than ex-ante optimal. Specifically, c is equal to what self 1 would borrow if she enjoyed the utility of consumption in that period, discounting the second installment of repayment by β. To see why this is the case, notice that the contract is designed so that self 0 expects to finish her repayment obligations in period 1 ($\hat{r} = 0$). Hence, when deciding whether to participate, self 0 trades off c with $k(\hat{q})$. But from the firm's perspective, $k(\hat{q})$ is just the highest actual total cost of repayment that can be imposed on self 1 so that she is still willing to change the installment plan. This means that the tradeoff determining the profit-maximizing level of borrowing is between c and *self 1's* cost of repayment, which discounts the second installment by β.

It is worth emphasizing that all the above holds for *any* $\hat{\beta} > \beta$—even for an arbitrarily small amount of naivete. Borrowers with $\hat{\beta} \approx \beta$ know that they will be much more impatient in the future than they would now like, and in principle have great demand for commitment. Nevertheless, they end up getting none in the market, and borrow too much to boot. There is therefore a discontinuity at full sophistication: fully sophisticated consumers borrow the right amount and repay according to self-0 preferences, but all other types of consumers borrow too much and repay according to period-1 preferences. In this sense, the equilibrium non-linear contract targets and exaggerates even an arbitrarily small amount of naivete, inducing the consumer to make a non-trivial mistake.

As the above intuition makes clear, non-sophisticated consumers have low welfare in the credit market because firms can offer contracts with discontinuous jumps between what the consumer expects to choose and what she ends up choosing. We now show that unless the consumer is very naive, there is a simple intervention that takes care of this problem:

[11] The prediction regarding the amount of borrowing contrasts in an interesting way with predictions of hyperbolic discounting in standard consumption-savings problems such as Laibson (1997). In those problems, as with linear contracts in our setting, whether more naive decisionmakers borrow more depends on the per-period utility function. In our setting, non-sophisticated consumers borrow more than sophisticated ones for any $k(\cdot)$.

Proposition 3. *A sophisticated consumer ($\hat{\beta} = \beta$) is equally well off in the restricted long-term market and the unrestricted long-term market. If a non-sophisticated consumer ($\hat{\beta} > \beta$) is sufficiently sophisticated ($\hat{\beta}$ is sufficiently close to β), she is strictly better off in the restricted long-term market than in the unrestricted long-term market.*

For sophisticated consumers, the restricted and unrestricted markets work equally well—they both generate the highest possible level of utility. By counteracting her tendency for immediate gratification as given by β, a contract with an interest rate of $R = 1/\beta$ aligns self 1's behavior with the borrower's long-run welfare. And since sophisticated consumers understand their own behavior perfectly, it is profit-maximizing to offer such a contract to them. Hence, what matters for sophisticated borrowers is that they can sign a long-term contract in period 0, and it is not crucial for those contracts to be unrestricted.

More interestingly, restricting contracts to have a linear structure prevents firms from fooling borrowers into discretely mispredicting their behavior, and hence raises the welfare of non-sophisticated but not-too-naive borrowers. For any interest rate R, a slightly naive borrower mispredicts her future behavior only by a small amount, and leads her to make only a small mistake in how much she borrows. This means that her behavior is very close to a sophisticated borrower's behavior, so that she will get an interest rate very close to that offered to a sophisticated borrower. As a result, her utility is very close to optimal.

Our formal result holds only for sufficiently sophisticated borrowers because if a borrower is very naive, even a linear contract can lead her to severely underpredict how much she will be willing to pay back in period 1. To exploit this misprediction, a firm may set such a high interest rate that the consumer's welfare falls below that in the unrestricted market . While technically Proposition 3 is stated as a limit result, it is clear that for this to happen for most utility functions $k(\cdot)$, the consumer has to be quite naive. In addition, if a social planner is worried that some consumers may indeed be very naive, a restricted market with an interest-rate cap will often still be better than an unrestricted market.

While our main interest is in actual behavior, the consumer's anticipated repayment schedule—the decoy option—is also intriguing. The firm attracts borrowers with an inefficiently "virtuous" option: it asks consumers to carry out all repayment in period 1, even though the marginal cost of repaying a little bit in period 2 could be zero. Intuitively, because the decoy repayment terms are never implemented, the firm is not motivated to design them efficiently. Instead, its goal in designing the decoy option is to extract as much profit as possible from the consumer's misunderstanding of her future behavior. Since the consumer overestimates her future self-control, the contract that is best at exploiting her misperception is one that asks her to exercise the most self-control.

3.2 The Role of Exclusivity

Our analysis in this section relies on the assumption that firms and consumers can write exclusive contracts in period 0, preventing a consumer from borrowing from another firm in period 1. In our three-period setting, if consumers could borrow externally in period 1, they could never be fooled into repaying more than they expected, because they could simply finance part of their period-1 repayment from another loan. Nevertheless, we argue that—especially in a more realistic, longer-horizon setting—the kinds of contracts we have derived are for several reasons likely to survive and determine outcomes and welfare in the market.

While we have modeled contracts as fully exclusive, a sufficient condition for our results is that the original firm can introduce a fee—such as the hefty refinancing penalties in predatory mortgages emphasized by Engel and McCoy (2002)—for changing repayment terms or contracting with another firm in period 1. These contracts resemble the type of contracts we have derived above, so non-sophisticated consumers will be willing to sign them. And since the original firm collects the fee and profits from the consumer's inclination to change the repayment schedule, our results survive essentially unchanged.

Even if there is no explicit penalty for contracting with another firm, borrowers' time inconsistency may prevent them from doing so. Looking for advantageous refinancing possibilities for an existing loan is (for most people) an unpleasant activity, and putting it off by a few days or weeks has a very small cost. As shown by O'Donoghue and Rabin (2001), a partially naive time-inconsistent decisionmaker often procrastinates indefinitely on such a task, believing that she will do it at the next opportunity.[12] Consistent with this view, Shui and Ausubel (2004) document that credit-card consumers rarely take advantage of opportunities to lower interest payments on existing debt.

Finally, even if a borrower refinances a loan, she is likely to do so using a long-term contract of the type we have identified, and if she keeps doing this, she eventually runs out of refinancing options. Indeed, Engel and McCoy (2002) document that predatory mortgages are often refinanced with similarly structured loans, until eventually the borrower cannot obtain any more credit and has to repay according to her existing contract. In the same vein, credit-card balance-transfer deals and teaser rates draw consumers into contracts similar to those they had before, and when the borrower reaches a stage where she can no longer roll over her debt, she has to pay according to the existing contract terms.

4 Non-Linear Contracting with Unknown Types

The insights of the previous section may appear to rely on the assumption that firms know β and $\hat{\beta}$—and use this information to fool even slightly naive borrowers. This section investigates market outcomes when either β, or both β and $\hat{\beta}$, are unobservable to firms. Somewhat surprisingly, our key intuitions survive: sophisticated consumers—although they are asked to repay their debt inefficiently quickly—benefit from the availability of long-term contracts, while all others follow a discontinuously different repayment schedule and have discontinuously lower welfare as a result. We also show that when different types of borrowers cannot be distinguished, non-sophisticated borrowers cross-subsidize sophisticated ones, and sophisticated borrowers "discipline" firms by preventing very misleading contracts.

Our analysis proceeds from a simple observation: the two-dimensional heterogeneity firms face in our most general setting is unusual because the two dimensions of private information do not affect consumer behavior at the same time. In particular, period-0 behavior is determined solely by $\hat{\beta}$, and (given the available options) period-1 behavior is determined solely by β. This suggests that

[12] Ironically, it might be that only sophisticated consumers look for outside financing, even though they are the only ones for whom an exclusive contract is beneficial.

consumers may separate according to $\hat{\beta}$ in period 0, and then according to β in period 1. Following this reasoning, we begin by investigating outcomes when firms know $\hat{\beta}$ but not β, and then identify conditions under which the same contracts endogenously separate consumers according to $\hat{\beta}$ in period 0. While in the text we allow only two types of β for any given $\hat{\beta}$, we comment on how our results would be affected with more general distributions, and in the appendix we consider the general case formally.

4.1 Known $\hat{\beta}$, Unknown β

Suppose that a borrower's $\hat{\beta}$ is known, and there are two possible β's: $\beta_1 < \hat{\beta}$ with probability p_1 and $\beta_2 = \hat{\beta}$ with probability p_2. That is, in a population of borrowers sharing the same $\hat{\beta}$, there is a proportion p_1 of non-sophisticated borrowers and a proportion p_2 of sophisticated borrowers. Our first observation is that because both types of customers have the same beliefs in period 0, they will accept the same contract. With this in mind, we apply Lemma 1 and set up an individual firm's problem as choosing a type-independent consumption c and a menu of type-dependent repayment options $(q_1, r_1), (q_2, r_2)$ subject to participation, incentive, and perceived-choice constraints. As in textbook models of screening (e.g. Bolton and Dewatripont 2005, Chapter 2), the incentive constraints hold if and only if (i) type 1 prefers the repayment schedule (q_1, r_1) to (q_2, r_2); and (ii) the monotonicity condition $q_2 \geq q_1$ is satisfied. We solve a relaxed problem with the constraint (i), and since the solution to the relaxed problem will satisfy (ii), incentive compatibility will be satisfied. Unlike in a standard screening problem, however, both types initially believe that they are type β_2, so there is only one participation constraint. For the same reason, the perceived-choice constraints are identical to the sophisticated type's incentive constraint, so we do not state them separately.

Given the above considerations, the firm's problem is

$$\max_{c,q_1,r_1,q_2,r_2} \quad p_1(q_1+r_1)+p_2(q_2+r_2)-c$$

$$\text{s.t.} \quad c-k(q_2)-k(r_2) \geq \underline{u}, \tag{PC}$$

$$-k(q_1)-\beta_1 k(r_1) \geq -k(q_2)-\beta_1 k(r_2). \tag{IC}$$

In the optimal solution, IC binds; otherwise, the firm could increase q_1 without violating IC, increasing profits. In addition, PC binds; otherwise, the firm could increase q_2 without violating PC and relaxing IC, again increasing profits. From the binding constraints, we can express $k(q_1)$ and $k(q_2)$:

$$\begin{aligned} k(q_2) &= c-k(r_2)-\underline{u}, \\ k(q_1) &= k(q_2)+\beta_1(k(r_2)-k(r_1)). \end{aligned}$$

Inverting these functions and plugging them into the principal's objective function yields:

Proposition 4 (Period-1 Screening)**.** *Suppose $\hat{\beta}$ is known, and β takes one of two values, $\beta_1 < \hat{\beta}$ or $\beta_2 = \hat{\beta}$, with probabilities p_1 and p_2, respectively. In a competitive equilibrium, the installment plans (q_1, r_1) and (q_2, r_2) satisfy*

$$k'(q_2)-k'(r_2) = \frac{1-\beta_1}{\beta_1}\cdot\frac{k'(r_2)}{k'(r_1)}\cdot\frac{p_1}{p_2}, \tag{1}$$

$$k'(q_1)-\beta_1 k'(r_1) = 0. \tag{2}$$

By Equation (1), the sophisticated type's repayment schedule calls for a first installment that is too high even from the long-term perspective of period 0. When considering whether to allocate more of the sophisticated borrower's repayment to period 2, the firm faces a trade-off. On the one hand, this increases the sophisticated borrower's period-0 utility, increasing the profits from selling to her. On the other hand, since a type-1 borrower does not like to repay early, the same adjustment benefits her even more, so that she is willing to pay less to delay repayment. This tradeoff is similar to that in standard screening problems between increasing efficiency for the less profitable type and decreasing the information rent paid to the more profitable type. In our model, however, the relevant preferences in this tradeoff exist at different times. Since a sophisticated borrower sticks to her period-0 repayment plan, the payment the firm can extract from her depends only on period-0 preferences, so that this side of the tradeoff takes the period-0 perspective. But since a non-sophisticated borrower abandons her desired repayment plan, the payment the firm can extract from her depends on period-1 preferences, so that this side of the tradeoff takes the period-1 perspective. As a result, the sophisticated borrower's schedule is distorted, from the *period-0* perspective, in the opposite direction than the non-sophisticated borrower's *period-1* preferences.

Equation (2) is a version of the "no-distortion-at-the-top" result common to many screening problems. Since type 1 has the highest demand for delaying repayment—so that there is no reason to distort her contract to lower the rents of more profitable types—she repays in an efficient way from the perspective of period 1. The difference between the sophisticated and non-sophisticated borrowers' first-order conditions leads to a generalization to unknown types of our insight above that there is a discontinuity in outcomes at full sophistication. In a standard screening model without time inconsistency of preferences, as one type's preference parameter approaches the other type's, the two bundles in the optimal contract converge. As $\beta_1 \nearrow \beta_2$ in our model, however, q_1 approaches a number strictly smaller than does q_2. In other words, a non-sophisticated borrower repays in a discontinuously different way from a sophisticated borrower, even if she is arbitrarily close to sophisticated. Furthermore, her repayment schedule caters fully to period-1 preferences. Obviously, this also means that she is discontinuously worse off than a sophisticated borrower.

It is worth emphasizing that the discontinuity in behavior and welfare in our model results from a combination of a misprediction of preferences *and* time inconsistency. Suppose the long-run preferences in our model discount period 2 by β, so that the consumer may mispredict her preferences, but those preferences are time consistent. Then, to exploit any misprediction of preferences, the optimal contract still induces unexpected switching. But if the misprediction is small, period-0 and period-1 preferences are close, so that there is no jump between the decoy option and the chosen option.

The above properties of the competitive-equilibrium contract arguably closely resemble some puzzling features of real-life credit arrangements. Loaded with cash-back bonuses, free rental-car insurance, a grace period, and other perks, the typical credit-card deal is extremely favorable—so long as the consumer does not revolve any debt on her card. If she leaves even $1 unpaid every month, she is charged interest on *all* purchases, and all of a sudden credit-card use becomes quite expensive. From a firm's cost perspective, it is unclear why a credit contract should require repayment in full within such a short time horizon, and be so much more expensive otherwise. But a favorable basic deal involving very fast repayment, and a discontinuity afterwards, is exactly what our model predicts. Similarly, in-store financing deals often involve no interest for a few months, but if a consumer does not repay fully within the allotted time, she is charged interest from the time of purchase on the entire purchase. Once again, the pattern of a good deal involving fast repayment along with a discontinuity afterwards is exactly what our model predicts. Finally, subprime mortgages often seem to have harsh repayment terms, and refinancing the mortgage carries large penalties.

These patterns are also consistent with a model of "shrouded attributes" in the sense of Gabaix and Laibson (2006): if some consumers are simply not aware of how grace periods or mortgages work, firms have an incentive to introduce hidden costs that take advantage of their misunderstanding. It seems that many consumers are indeed unaware of how exactly credit works, and this certainly contributes to their behavior in credit markets. We view our model as providing a complementary analysis. Whereas Gabaix and Laibson (2006) assume essentially exogenously what consumers ignore when making purchase decisions (e.g. that they ignore a printer's cartridge costs), we derive this misprediction from a general model of consumer preferences interacting with profit-maximizing firms in an unrestricted market. This allows us to derive more specific predictions regarding credit contracts. For example, we predict that in decisions implicating time inconsistency, but not in other decisions, contracts will have a discontinuous structure. It is not clear why a model of shrouded attributes would necessarily have such a feature.

4.2 Welfare

In this section, we consider the welfare of borrowers in the above market with $\hat{\beta}$ known and β unknown. Our main interest is in comparing welfare to that in the two restricted markets:

Proposition 5. *Suppose $\hat{\beta}$ is known, and β takes one of two values, $\beta_1 < \hat{\beta}$ or $\beta_2 = \hat{\beta}$, with probabilities p_1 and p_2, respectively. Consumers strictly prefer accepting the contract in the unrestricted market over participating in the short-term market or the long-term restricted market, and a sophisticated consumer is indeed better off. If Assumption 1 holds, the non-sophisticated consumer has strictly lower welfare than in the short-term market. If the non-sophisticated consumer is sufficiently sophisticated (β_1 is sufficiently close to $\hat{\beta}$), her welfare, as well as total social welfare, is greater in the long-term restricted market than in the long-term unrestricted market.*

A consumer strictly prefers the long-term unrestricted market over the long-term restricted or the short-term one; otherwise, a firm could offer a contract replicating the preferred market. And since a sophisticated consumer accurately predicts her utility, she is actually strictly better off. In contrast, the discontinuity in welfare between sophisticated and non-sophisticated consumers is large enough that the latter are worse off than in the short-term market. Hence, even when β is unobserved, unrestricted contracts leave non-sophisticated consumers with rather low welfare.

Perhaps most importantly, restricting the long-term market to be linear raises the welfare of non-sophisticated consumers as well as total welfare in the population. The basic reason is the same as when β is known: linear markets prevent non-sophisticated but not-too-naive borrowers from drastically mispredicting their behavior. In the current setting, however, the long-term restricted market does not Pareto-dominate the unrestricted one—it makes sophisticated borrowers worse off. Intuitively, because non-sophisticated borrowers are more profitable, in a competitive equilibrium it must be that firms make money on non-sophisticated borrowers and lose money on sophisticated borrowers. This cross-subsidy benefits sophisticated borrowers. Nevertheless, the restricted market is still socially optimal since it decreases the distortion in the consumption and repayment schedules to both types.

In the rest of this section, we analyze how welfare depends on the proportion of non-sophisticated types in the population. The welfare of sophisticated borrowers is strictly increasing in p_1:

Proposition 6. *Suppose $\hat{\beta}$ is known, and β takes one of two values, $\beta_1 < \hat{\beta}$ or $\beta_2 = \hat{\beta}$, with probabilities p_1 and p_2, respectively. The sophisticated type's utility in the competitive equilibrium is strictly increasing in p_1.*

The intuition for Proposition 6 is easiest to see based on the cross-subsidies above. Since firms make money on non-sophisticated borrowers, the immediate effect of an increase in p_1 is an increase in profits. With competition, however, this leads firms to compete more fiercely for consumers, increasing the welfare of sophisticated borrowers.

Based on the cross-subsidy effect, it is natural to conjecture that an increase in p_1 also benefits non-sophisticated borrowers by lowering the number of borrowers to subsidize and raising the number of borrowers to subsidize them. This cross-subsidy effect was emphasized by Gabaix and Laibson (2006). In our model, however, there is a force acting in the opposite direction. As p_1 increases, the mistake non-sophisticated consumers make in estimating their utility increases:

Proposition 7 (The Discipline Effect)**.** *Suppose $\hat{\beta}$ is known, and β takes one of two values, $\beta_1 < \hat{\beta}$ or $\beta_2 = \hat{\beta}$, with probabilities p_1 and p_2, respectively. Then, $[c - k(q_2) - k(r_2)] - [c - k(q_1) - k(r_1)]$ is increasing in p_1.*

Intuitively, a firm can make more money from sophisticated consumers by offering them a more efficient decoy option, but it can make more money from non-sophisticated consumers by increasing the extent to which they misestimate their desire to delay payment and hence their utility. As the number of non-sophisticated consumers increases, the latter consideration becomes more important, hurting non-sophisticated consumers. In this sense, sophisticated consumers provide market "discipline" by forcing firms to offer a more reasonable decoy option.[13] As a result, the effect of an increase in p_1 on the welfare of non-sophisticated borrowers is ambiguous.

As we show in the appendix, the main conclusions we have derived in a framework with two types of β given $\hat{\beta}$ hold for more general distributions. The main additional issue is that when there are multiple non-sophisticated types in a population of borrowers with given beliefs, the

[13] The discipline effect is reminiscent of the positive externality informed consumers exert on uninformed consumers in the pricing models of Salop and Stiglitz (1977) and Wolinsky (1983). Since firms lose informed consumers if they increase premiums, these consumers serve to decrease prices and price dispersion. In contrast to our model, however, in these theories uninformed consumers have a negative externality on informed consumers.

optimal contract screens between them in period 1. This screening redistributes money from lower-β to higher-β types, and as a side benefit provides an incentive to repay earlier. Specifically, to lower the information rents paid to the most time-inconsistent borrowers, less time-inconsistent types repay more quickly than is optimal from a period-1 perspective. Nevertheless, there is still a tendency for non-sophisticated consumers to have discontinuously different behavior and welfare than sophisticated consumers, and if non-sophisticated consumers are not too naive, the long-term restricted market yields higher social welfare than the unrestricted market.

4.3 Unknown β and $\hat{\beta}$

We now build on the insights above to identify the competitive equilibrium when consumers differ in both β and $\hat{\beta}$. We provide a condition under which consumers self-select according their perceived preferences $\hat{\beta}$ in period 0, and then according to their true preferences β in period 1. This means that even with two-dimensional asymmetric information, non-sophisticated consumers, even if very close to sophisticated, endogenously select credit contracts in which—to their detriment—they end up changing their mind about repayment.

Throughout this section, suppose that the possible β's are $\beta_1 < \beta_2 < \cdots < \beta_I$, and $\hat{\beta} \in \{\beta_2, \ldots \beta_I\}$. For any given $\hat{\beta} = \beta_i$, the borrower has $\beta = \beta_i$ with probability p_i and $\beta = \beta_{i-1}$ with probability $1 - p_i$. In the appendix, we allow for more general distributions of β given $\hat{\beta}$. Let $\underline{u}(\beta_i, \beta_{i-1}, p_i)$ be a borrower's perceived utility in the competitive equilibrium when $\hat{\beta} = \beta_i$ is observable, with probability p_i the borrower is sophisticated, and with probability $(1 - p_i)$ she is type β_{i-1}. Our key condition is the following:

Condition 1. $\underline{u}(\beta_i, \beta_{i-1}, p_i)$ is increasing in β_i.

Condition 1 states that if $\hat{\beta}$ was observable, the perceived utility from the equilibrium contract would be increasing in $\hat{\beta}$. While this is an endogenous condition, it is intuitively plausible. The condition requires roughly that borrowers who are more optimistic about their future behavior tend to be more naive about it. Since firms compete more fiercely for such profitable consumers, they drive up the perceived attractiveness of the deal. In the current setting with two types of

β for each $\hat{\beta}$, we require that consumers who believe themselves to be less time-inconsistent are non-sophisticated with sufficiently higher probability.[14]

We now argue that under Condition 1, there is a competitive equilibrium in which firms offer the same contracts we have derived in Section 4.1. The crucial part is that from such a set of contracts, consumers self-select according to $\hat{\beta}$ in period 0; then, since firms cannot do better even when they know $\hat{\beta}$, they certainly cannot do better when they do not know $\hat{\beta}$. There are two parts to our self-selection argument. First, since a borrower of type $\hat{\beta}$ expects to follow the decoy repayment option in a contract intended for any $\hat{\beta}' \leq \hat{\beta}$, among these contracts she prefers the one intended for her simply because (by Condition 1) it gives her the highest perceived period-0 utility. Second, while from a period-0 perspective the borrower prefers the decoy option in the contract for $\hat{\beta}' > \hat{\beta}$ to the decoy option in the contract for her own type, she also believes that she will switch away from this option ex post. Once she takes this into account, the period-0 utility from the contract designed for $\hat{\beta}' > \hat{\beta}$ is lower. To see this last point, suppose by contradiction that a type $\hat{\beta}$ preferred to select the contract designed for $\hat{\beta}' > \beta$. Then, the contract for $\hat{\beta}$ must be suboptimal when $\hat{\beta}$ is known. In particular, the contract designed for $\hat{\beta}'$ both attracts $\hat{\beta}$ types and induces all of them to choose the non-sophisticated repayment option. Because this is the more profitable repayment option from the firm's perspective, the contract thus yields strictly positive profits.

Intuitively, the credit contract intended for a borrower who is more optimistic about her future behavior offers a better deal if the borrower can stick to the most favorable repayment schedule, but requires greater self-control to stick to that schedule. Hence, a consumer takes the most favorable credit contract with which she believes she can still repay fast. If the consumer is not perfectly

[14] It is easy to see that Condition 1 is satisfied for some distributions of types. For example, suppose the possible β's are $\beta_1, \beta_2, \beta_3$, and p_3 is fixed with $0 < p_3 < 1$. Then, as $p_2 \to 1$, Condition 1 holds. To see this, suppose the opposite was true, i.e. $\underline{u}(\beta_2, \beta_1, 1) \geq \underline{u}(\beta_3, \beta_2, p_3)$. Recall from above, that as $p_2 \to 1$, the repayment option selected by sophisticated consumers is determined by $k'(q_2) = k'(r_2) = 1$ and the zero profit condition $c = q_2 + r_2$. Since $\underline{u}(\beta_2, \beta_1, 1) \geq \underline{u}(\beta_3, \beta_2, p_3)$ the contract designed for $\hat{\beta} = \beta_2$ would satisfy the participation condition also in the market in which $\hat{\beta} = \beta_3$. Furthermore, when offering the contract designed for $\hat{\beta} = \beta_2$ to $\hat{\beta} = \beta_3$ a firm would break even. But for the given outside option, firms can do strictly better by offering the optimal contract we solved for above. Thus, in case $\underline{u}(\beta_2, \beta_1, 1) \geq \underline{u}(\beta_3, \beta_2, p_3)$ firms make positive profits in the market in which $\hat{\beta} = \beta_3$, yielding the desired contradiction. In the Appendix, where we consider an extended screening model in which an agent with beliefs $\hat{\beta} = \beta_i$ may be any period-1 type $\beta_j \leq \beta_i$, we provide an entire class of distributions of types $(\hat{\beta}, \beta)$ that satisfy Condition 1.

sophisticated, however, this is a contract with which she does not repay fast.

To illustrate the logic of self-selection through an example, consider a consumer looking to buy a TV on sale financed using store credit that does not accrue interest for six months. The nicer the TV, the sweeter is the deal both because the sale is steeper and because the six-month interest-free period is more valuable. At the same time, it is more difficult to pay back a larger loan in six months. Hence, the consumer chooses a TV which she believes she will just finish paying off in time. But if she is even slightly naive, this TV will be too nice, and she will fail to pay it off.

In fact, the above competitive equilibrium is the unique one:[15]

Proposition 8 (Period-0 Screening)**.** *Suppose Condition 1 holds. Then, in the unique equilibrium with $\hat{\beta}$ unobserved by firms, each consumer accepts the same contract as when $\hat{\beta}$ is observed by firms.*

This conclusion is a general version of our insight that long-term contracts will not help borrowers achieve self-control. Even when firms observe neither the consumer's preferences nor her degree of sophistication, all unsophisticated consumers—no matter how close they are to sophistication—endogenously select a contract that takes no advantage of the commitment technology available. Self-selection therefore ensures that the discontinuity we have identified when the firm knows both β and $\hat{\beta}$ does not depend on such knowledge.

We conclude the paper by discussing how the welfare of an infinitesimally small consumer with a given β changes if she becomes more sophisticated. Clearly, this question requires a model with heterogeneity in both β and $\hat{\beta}$. With the current distribution of types, when a non-sophisticated person becomes more sophisticated, she becomes fully sophisticated. By simple revealed-preference considerations, this makes her better off. One might more generally conjecture that an increase in sophistication—a decrease in $\hat{\beta}$—increases a borrower's utility from a period-0 perspective. Using our more general model, Appendix A shows through an example that this is false, and in the example is false in a dramatic way: the borrower moves from completely naive to near sophisticated, yet this makes her worse off. Intuitively, so long as a borrower does not become *fully* sophisticated,

[15] The intuition for why the competitive equilibrium must involve full sorting of $\hat{\beta}$'s is the following. If different consumers accepted the same contract in period 0, a competing firm could "steal" and make money on the most valuable type in that pool by offering the optimal screening contract for that type.

she switches away from her preferred repayment schedule to one that is closer to self 1's preferences. Hence, becoming more sophisticated does not help in achieving self-control in repayment. In addition, if she is in a pool of borrowers who are less optimistic (and hence less likely to be naive) about their future preferences, firms give her a worse deal up front, lowering her utility.

5 Conclusion

While the kinds of exploitative contracts we have derived capture some salient features of real-world credit contracts, there are other seemingly important questions we have not addressed. Perhaps most importantly, consumers may learn about their preferences from their own behavior and that of the firms, and they often seem to have a generic skepticism regarding contracts even if they do not know how exactly the contract is looking to exploit them.[16] Since our model (and most models of naivete with which we are familiar) starts from exogenously given beliefs, it cannot easily accommodate such learning.

References

Agarwal, S., J. C. Driscoll, X. Gaibaix, and D. Laibson (2007): "The Age of Reason: Financial Decisions Over the Lifecycle," MIT Department of Economics Working Paper No. 07-11.

Angeletos, G.-M., D. I. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001): "The Hyperbolic Consumption Model: Calibration, Simulation, and EmpiricalEvaluation," *Journal of Economic Perspectives*, 15(3), 47–68.

Asheim, G. B. (2007): "Procrastination, Partial Naivete, and Behavioral Welfare Analysis," Working Paper, University of Oslo.

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

——— (1999): "Adverse Selection in the Credit Card Market," Working Paper, University of Maryland.

[16] While the issue of whether and how people learn about their time inconsistency has not been carefully investigated, evidence by Agarwal, Driscoll, Gaibaix, and Laibson (2007) suggests more broadly that consumers learn to make fewer mistakes in financial decisions, at least until their cognitive abilities start to significantly decline in their 50's.

Bernheim, B. D., and A. Rangel (2004): “Addiction and Cue-Triggered Decision Processes,” *American Economic Review*, 94(5), 1558–1590.

Bertaut, C. C., and M. Haliassos (2002): “Debt Revolvers for Self Control,” Mimeo.

Bolton, P., and M. Dewatripont (2005): *Contract Theory*. MIT Press.

DellaVigna, S., and U. Malmendier (2004): “Contract Design and Self-Control: Theory and Evidence,” *Quarterly Journal of Economics*, 119(2), 353–402.

——— (2006): “Paying Not to Go to the Gym,” *American Economic Review*, 96(3), 694–719.

DellaVigna, S., and M. D. Paserman (2005): “Job Search and Impatience,” *Journal of Labor Economics*, 23(3), 527–588.

Durkin, T. A. (2000): “Credit Cards: Use and Consumer Attitudes, 1970-2000,” Federal Reserve Bulletin.

Eliaz, K., and R. Spiegler (2006): “Contracting with Diversely Naïve Agents,” *Review of Economic Studies*, 73(3), 689–714.

Engel, K. C., and P. A. McCoy (2002): “A Tale of Three Markets: The Law and Economics of Predatory Lending,” *Texas Law Review*, 80(6), 1255–1381.

Fang, H., and D. Silverman (2007): “Time-inconsistency and Welfare Program Participation: Evidence from theNLSY,” Working Paper, University of Michigan.

Gabaix, X., and D. Laibson (2006): “Shrouded Attributes, Consumer Myopia, and Information Suppression in CompetitiveMarkets,” *Quarterly Journal of Economics*, 121(2), 505–540.

Gruber, J., and B. Kőszegi (2004): “A Theory of Government Regulation of Addictive Bads: Optimal Tax Levelsand Tax Incidence for Cigarette Taxation,” *Journal of Public Economics*, 88(9-10), 1959–1987.

Heidhues, P., and B. Kőszegi (2007): “Futile Attempts at Self-Control,” Working Paper in Progress.

Kőszegi, B. (2005): “On the Feasibility of Market Solutions to Self-Control Problems,” *Swedish Economic Policy Review*, 12(2), 71–94.

Laibson, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 112(2), 443–477.

Laibson, D. I., A. Repetto, and J. Tobacman (1998): “Self-Control and Saving for Retirement,” in *Brookings Papers on Economic Activity*, ed. by W. C. Brainard, and G. L. Perry, vol. 1, pp. 91–196. The Brookings Institution.

——— (2007): “Estimating Discount Functions with Consumption Choices over the Lifecycle,” Working Paper, Harvard University.

Livshits, I., J. MacGee, and M. Tertilt (2006): “Accounting for the Rise in Consumer Bankruptcies,” Stanford University Working Paper.

——— (2007): “Consumer Bankruptcy: A Fresh Start,” *American Economic Review*, 97(1), 402–418, American Economic Review.

O’Donoghue, E., and M. Rabin (2006): “Optimal Sin Taxes,” *Journal of Public Economics*, 90(10-11), 1825–1849, Working Paper, University of California, Berkeley.

O’Donoghue, T., and M. Rabin (2001): “Choice and Procrastination,” *Quarterly Journal of Economics*, 116(1), 121–160.

Oster, S., and F. Scott-Morton (2005): “Behavioral Biases Meet the Market: the Case of Magazine Subscription Prices,” *Berkeley Electronic Journals in Economic Analysis & Policy — Advances*, 5(1).

Paserman, M. D. (2006): “Job Search and Hyperbolic Discounting: Structural Estimation and PolicyEvaluation,” Working Paper.

Salop, S., and J. Stiglitz (1977): “Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion,” *Review of Economic Studies*, 44(3), 493–510.

Shui, H., and L. M. Ausubel (2004): “Time Inconsistency in the Credit Card Market,” Available at http://ssrn.com/abstract=586622.

Strotz, R. H. (1956): “Myopia and Inconsistency in Dynamic Utility Maximization,” *Review of Economic Studies*, 23, 165–180.

Wolinsky, A. (1983): “Prices as Signals of Product Quality,” *Review of Economic Studies*, 50(4), 647–658.