Competition over Time-Inconsistent Consumers

DANIEL GOTTLIEB
Department of Economics, Massachusetts Institute of Technology

Abstract

How do firms respond to consumers' time inconsistency? This paper studies the optimal design of nonexclusive contracts under competition. It shows that nonexclusivity creates a stark asymmetry between immediate-costs goods and immediate-rewards goods. For immediate-cost goods nonexclusivity does not affect the equilibrium and, when consumers are sophisticated, the efficient allocation is achieved. When consumers are partially naive, the optimal sales tax may be either positive or negative and depends on parameters that are hard to estimate. In the case of immediate-rewards goods, however, the equilibrium features marginal-cost pricing and is always inefficient. Moreover, the optimal tax does not depend on the consumers' degree of naiveté and is a function of parameters that are easy to assess.

1. Introduction

A significant amount of evidence suggests that, in some markets, consumers are not time consistent. When making intertemporal decisions, they usually give higher relative weights to an earlier future period as it gets closer. Consumers also tend to overestimate their degree of time consistency.

DellaVigna and Malmendier (2006), for example, analyze health clubs that offer both monthly contracts and 10-visit passes. Most consumers take the monthly contracts even though they would spend less if they bought a 10-visit pass. Shui and Ausubel (2005) consider a randomized experiment in the credit card industry. Consumers are offered both a contract with a lower

Daniel Gottlieb, MIT Department of Economics, E52-391, 50 Memorial Drive, Cambridge, MA 02142 (daniel_g@mit.edu).

I am grateful to Mushe Cohen, Stefano DellaVigna, Glenn Ellison, Xavier Gabaix, Jonathan Gruber, Botond Köszegi, David Laibson, Johannes Spinnewijn, and especially Jean Tirole for comments. Remaining errors are mine.

Received May 21, 2007; Accepted March 6, 2008.

Journal of Public Economic Theory, 10 (4), 2008, pp. 673–684.

interest rate and a shorter duration and one with a higher interest rate but a longer duration. Although most consumers take the contract with a shorter duration, they would pay less if they chose the one with longer duration since their choice induces them to continue to borrow on the credit card.[1]

How do firms respond to time inconsistency among consumers? DellaVigna and Malmendier (2004) have shown that firms set prices above marginal cost for goods with immediate rewards and deferred costs (*leisure goods*) and below marginal cost for goods with immediate costs and deferred rewards (*investment goods*) both in the monopolistic and in the competitive case. This result follows from the fact that commitment devices are valued by consumers and, therefore, firms are able to benefit from providing them. Furthermore, as long as consumers are aware of their being time inconsistent, the equilibrium is efficient.

Although DellaVigna and Malmendier (2004) show that their model is consistent both with evidence from investment goods such as the health club industry and the vacation time-sharing industry as well as with evidence from leisure goods such as the credit card and the mobile phone industries, it is not consistent with observations from the tobacco, alcohol, and unhealthy food industries (which are probably among the most compelling examples of markets with time-inconsistent consumers). Since these are leisure goods, the model implies that consumers would receive a lump-sum transfer and pay a usage price higher than the marginal cost. However, we do not observe such lump-sum transfers to consumers in these three examples.

In this paper, I argue that an important component missing from the DellaVigna and Malmendier model is the fact that firms are typically unable to offer exclusive contracts. In some markets (such as the cigarettes, alcohol, and unhealthy foods markets), the firms' cost of enrolling a potential user and the consumers' cost of switching between different firms is usually very low and, therefore, contracts are nonexclusive. Then, if a firm offers a positive lump-sum transfer to consumers and charges high usage prices (as happens in the case of leisure goods with exclusivity), another firm may benefit from selling the good at a lower price after the contract has been signed. Hence, the impossibility to prevent a consumer from buying from another firm after the contract has been signed imposes an important restriction on leisure goods markets.

In equilibrium, there is a stark asymmetry between leisure and investment goods. While the equilibrium in the investment goods case is the same as in DellaVigna and Malmendier (2004) and there is no role for taxation when consumers are aware of their time inconsistency, in the leisure goods case

[1]Oster and Scott-Morton (2005) consider the relationship between newsstand and subscription prices of magazines. They show that magazines associated with future benefits ("investment magazines") have relatively higher subscription costs than those with immediate benefits ("leisure magazines"), indicating that firms may be exploiting the consumers' time inconsistency.

prices are equal to marginal cost and the equilibrium is always inefficient. The key idea is that time-inconsistent consumers face an ex post incentive to circumvent the prearranged commitment devices. In the case of leisure goods, firms are able to profit ex post by offering these contracts, which breaks down the equilibrium.[2] In the case of investment goods, they are not able to profit, and the equilibrium with commitment devices is maintained. This prediction seems to be consistent with evidence that markets for leisure goods with small costs of enrolling new users and small switching costs do not feature lump-sum transferences to consumers.[3]

An implication of the model is that, in the case of leisure goods, increased competition may decrease total surplus. If a monopolist could commit not to renegotiate the contract but firms could not prevent consumers from signing contracts with other firms, total surplus in leisure goods markets would be higher under monopoly than under competition.[4]

The efficient allocation in the case of leisure goods can be obtained by a sales tax that corrects for the "internalities" of consuming the good. It does not depend on the degree of naiveté and requires a relatively small amount of information from a regulator. On the other hand, the optimal tax in the case of investment goods depends on the consumers' naiveté (the tax is zero if they are sophisticated) and requires much more information. Therefore, my model suggests that there is a much larger role for government intervention in leisure goods markets than in investment goods markets.

This paper is related to a growing literature on behavioral industrial organization.[5] Gabaix and Laibson (2006) analyze information shrouding by firms serving both rational and myopic consumers. They show that, under some conditions, informational shrouding persists even in competitive environments. Spiegler (2006) assumes that consumers use a sampling procedure because they have a limited ability to evaluate complex objects. In his model, increased competition leads to more obfuscation instead of the convergence of prices to marginal costs.[6]

---

[2]In that sense, my model is similar to the models of commitment not to renegotiate of Hart and Tirole (1988) and Dewatripont (1988). However, the issue here is not the incentive to renegotiate the contract but to accept another contract from a different firm.

[3]In an independent work, Köszegi (2005) has recently provided a similar suggestion for why we do not observe lump-sum transfers in some markets for leisure goods. He does not, however, explicitly model competition. He also does not study the welfare properties of the resulting equilibrium. Nevertheless, as in the model presented in this paper, he also obtains an asymmetry between markets for leisure and investment goods.

[4]This result relies on the assumption that consumers' preferences are homogeneous; therefore, there is no deadweight loss from monopoly. If consumers were heterogeneous, the result would depend on whether the deadweight loss from monopoly would overweight the gain from providing commitment.

[5]Ellison (2006) provides a survey of this literature.

[6]Eliaz and Spiegler (2006) consider a model where a monopolist faces consumers with limited ability to predict changes in their future tastes. Heidhues and Köszegi (2005) study the optimal contract offered by a monopolist facing loss-averse consumers.

The paper is also related to the literature on the optimal regulation of goods consumed by time-inconsistent agents. Gruber and Köszegi (2001) extended the Becker and Murphy's (1988) rational-addiction model to the case of quasi-hyperbolic consumers. O'Donoghue and Rabin (2003) studied the optimal taxation of unhealthy goods based on Ramsey's commodity taxation model. The main conclusion of these papers is that it is optimal to tax leisure goods and to subsidize investment goods.[7] The main reason for taxing leisure goods is the provision of "internalities": Taxes provide incentives for consumers to act according to their long-run preferences. In other words, taxing immediate-rewards goods provides a commitment device that prevents consumers from being tempted to behave differently from how their long-run selves would act.

In both papers, however, it is assumed that prices are set equal to the marginal cost in the absence of regulation. However, as shown by DellaVigna and Malmendier (2004), in the presence of time-inconsistent consumers, firms do not set prices equal to marginal cost. In this paper, I consider the optimal taxes when prices are set endogenously by competitive firms when contracts are nonexclusive.

2. The Model

The model is a competitive version of DellaVigna and Malmendier (2004). There are three periods. In the first period, consumers make a take-it-or-leave-it offer of a contract to firms. Because the good is indivisible, there is no loss of generality in assuming that a first-period contract consists of a two-part tariff (L, p), where L denotes a lump-sum price and p is a usage price. Both L and p are paid in period 2.

Consumption occurs in period 2. The good provides an immediate payoff (in period 2) of $-c$ and a delayed payoff (in period 3) of b. If $-c < 0 < b$, the good generates an initial cost and a delayed benefit. These goods are called *investment goods*. If $b < 0 < -c$, then the good generates an immediate benefit and a delayed cost. These are called *leisure goods*.

Consumers and firms have a common prior about F, the distribution of c. F is assumed to be a twice continuously differentiable distribution function with strictly positive density f on $\Theta \subset \mathbb{R}$. We take $\Theta = \mathbb{R}_+$ for the case of investment goods $(-c < 0 < b)$ and $\Theta = \mathbb{R}_-$ for the case of leisure goods $(b < 0 < -c)$.

In period 2, the realization of the immediate cost c is observed. Consumers make a take-it-or-leave-it offer of a contract to firms. A second-period

[7]Gruber and Köszegi have estimated the optimal cigarette taxes based on their model. They found that the provision of "internalities" due to hyperbolic discounting leads to an optimal tax of at least $1 per pack more than the traditional model.

contract is a price $\hat{p}$ contingent on consuming the good.[8] Then, they decide whether to consume the good.

If consumption occurred in period 2, consumers get a payoff with expected value b in period 3. Since there are no choices in period 3, it is irrelevant whether b is deterministic or stochastic.[9]

Therefore, the timing of the game is as follows:

1. Consumers offer a two-part tariff (L, p) to firms, where L denotes a lump-sum price and p is a usage price (both paid in period 2).
2. The immediate cost c is drawn from the distribution F. Consumers offer a price $\hat{p}$ contingent on consuming the good. Then, they decide whether to consume the good.
3. If the good was consumed in period 2, the consumer gets a payoff with expected value b.

Consumers have quasi-hyperbolic preferences:

$$U_t = u_t + \beta \sum_{s>t} \delta^{s-t} u_s,$$

where u_s is the instantaneous utility in period s and U_t is the discounted utility in period t. They are *exponential* (or time consistent) when their time-consistency parameter β is equal to 1. If β is less than 1, they are *hyperbolic* (or time inconsistent). A *partially naive* hyperbolic agent has true time-consistency parameter β, but believes that in the future she will behave like a hyperbolic agent with parameter $\hat{\beta} \in [\beta, 1]$. When $\hat{\beta} = \beta$, the agent is *sophisticated*.

Note that the discount rate that quasi-hyperbolic consumers use to discount the next period is lower than the discount rate with which they discount all subsequent future periods. Formally, at time t they impose a weight of $\beta\delta$ on period $t + 1$ and $\beta\delta^s$ on period $t + s$. Hence, the discount factor between periods t and $t + 1$ is $\beta\delta$ whereas the discount factor between $t + s$ and $t + s + 1$ is $\delta > \beta\delta$.

Each firm faces a cost of providing the good equal to $a > 0$, incurred in period 2 (which is when production occurs).[10] We assume that the firm has access to a credit market and faces a discount factor of δ.

Note that the offer of the two-part tariff is made under symmetric information. Therefore, a time-consistent consumer would choose marginal cost pricing and would extract all profits through the fixed fee (which, in this case,

[8]There is no loss of generality on assuming that the price if the good is not consumed is zero.

[9]Notice that delayed benefits b do not depend on the realization of c.

[10]We assume that there are no costs of signing a contract. Therefore, the firm faces no cost if a consumer signs a contract but does not consume the good. Our results still hold if these costs are small. On the other hand, assuming the presence of costs of signing a contract is equivalent to ruling out nonexclusive contracts if the costs are high enough.

would lead to $L = 0$, $p = a$). The equilibrium in the general case is obtained by backward induction.

If a consumer has not accepted a contract in the first period, her second period program (conditional on consuming the good) consists on choosing the price $\hat{p}$ that maximizes her utility subject to leaving nonnegative profits to the firms:

$$\max_{\hat{p}} \beta\delta b - \hat{p} - c$$
$$s.t. \quad \hat{p} \geq a.$$

The unique solution to this program is $\hat{p} = a$. If she has accepted a first period contract (L, p), her program (conditional on consuming the good) is

$$\max_{\hat{p}} \beta\delta b - \min\{\hat{p}, p\} - c$$
$$s.t. \quad \hat{p} \geq a.$$

The solution of this program is $\hat{p} = a$ if $p > a$ and $\hat{p} \geq a$ otherwise, and the consumer prefers the second period contract if $p > a$.

If a consumer accepts a first-period contract and consumes the good in the second period, she expects to obtain a benefit of $\hat{\beta}\delta b$ and faces a cost of $c + p$ in the second period. Hence, she expects to consume the good with probability $F(\hat{\beta}\delta b - p)$ and gets an expected discounted utility of $\beta\delta[\pi(p) - L]$, where $\pi(p) := \int_{-\infty}^{\hat{\beta}\delta b - p} (\delta b - p - c)\, dF(c)$ is the expected consumer surplus net of the usage price p.

The fact that the consumer prefers the second period contract if $p > a$ places an important restriction on the set of contracts that can be demanded in the first period. Since she will not consume using a contract (L, p) with $p > a$, the firm would get negative profits if it accepts $L \leq 0$. Therefore, we cannot have contracts with $p > a$ and $L \leq 0$.

Clearly, a consumer would never offer a first-period contract with $L \leq 0$ and $p > a$ since it would leave strictly positive profits to the firm (she could improve by offering $L' = 0$ and the same usage price). Hence, the possibility of offering contracts in the second period implies that first-period contracts must satisfy $p \leq a$.

Conversely, the consumer prefers to use a first-period contract whenever $p < a$. Hence, when $p < a$, the zero-profit condition for the firm becomes

$$L + F(\beta\delta b - p)(p - a) \geq 0.$$

Thus, any contract such that $p \leq a$ and $L + F(\beta\delta b - p)(p - a) \geq 0$ is accepted by the firms.

Therefore, the consumer's first-period program conditional on accepting a first-period contract is[11]

$$\max_{p,L} \beta\delta[\pi(p) - L]$$
$$s.t. \quad p \leq a$$
$$L + F(\beta\delta b - p)(p - a) \geq 0.$$

For the moment, ignore the $p \leq a$ constraint. Then, the first-order condition yields

$$p - a = -\delta b(1 - \hat{\beta})\frac{f(\hat{\beta}\delta b - p)}{f(\beta\delta b - p)} - \frac{F(\hat{\beta}\delta b - p) - F(\beta\delta b - p)}{f(\beta\delta b - p)}, \qquad (1)$$

which is the same expression as in DellaVigna and Malmendier (2004). The per-unit price p differs from the marginal cost a for two reasons. First, time-inconsistent consumers use prices as commitment devices as long as they are not fully naive (i.e., $\hat{\beta} < 1$). This is reflected by the term $-\delta b(1 - \hat{\beta})\frac{f(\hat{\beta}\delta b - p)}{f(\beta\delta b - p)}$. Second, by underestimating the probability of consuming the good, the consumer does not take into account the full cost of higher usage prices (captured by the second term).

In the case of investment goods, the term on the right-hand side of Equation (1) is negative. Therefore, the $p \leq a$ constraint does not bind and the equilibrium is exactly the same as the one in DellaVigna and Malmendier. The lump-sum price is then given by the zero-profit condition,

$$\begin{aligned} L &= -F(\beta\delta b - p)(p - a) \\ &= F(\beta\delta b - p)\left[\delta b(1 - \hat{\beta})\frac{f(\hat{\beta}\delta b - p)}{f(\beta\delta b - p)} + \frac{F(\hat{\beta}\delta b - p) - F(\beta\delta b - p)}{f(\beta\delta b - p)}\right]. \end{aligned}$$

In the case of leisure goods, the term is positive. Hence, the constraint binds and the solution is $p = a$. The zero-profit condition implies that the lump-sum price is $L = 0$. Therefore, *the equilibrium in the case of leisure goods involves marginal cost pricing*.

Note that, because leisure goods are priced at marginal cost, the model implies that prices are uninformative about the consumers' degree of time inconsistency when contracts are nonexclusive.

The model appears to be compatible with evidence from several markets. As DellaVigna and Malmendier (2004) argue, there is evidence that firms

[11] I assume that aggregate surplus is strictly concave for $p \leq a$. A sufficient condition is that $\frac{f'(c)}{f(c)}(\delta b - a - c) < 1$ for all $c \geq \beta\delta b - a$. This is satisfied, for example, when c is uniformly distributed.

Since the consumer can choose $L = 0$, $p = a$, there is no loss of generality in assuming that the first-period contract is accepted.

price investment goods such as health clubs and vacation time-sharing below marginal cost. On the other hand, leisure goods such as tobacco, alcohol, and unhealthy food do not feature ex ante lump-sum transfers to consumers and higher unit prices.

An important assumption of the model is that consumers are able to buy from other firms in the period that they consume the good. If contracts could only be signed in periods before consumption occurs, nonexclusivity would not lead to marginal cost pricing. In the credit card industry, for example, there is a time lag between applying for a credit card and being able to use it. Therefore, even though credit cards are clearly nonexclusive, usage prices are higher than marginal costs.

3. Welfare Analysis

This section studies the welfare properties of the equilibrium and characterizes the optimal taxes for leisure and investment goods. It is not entirely clear how to treat time inconsistency from a social welfare perspective. Should the short-run self perspective ($\beta < 1$) play a role in the welfare criterion?

Some papers have avoided this issue by applying a Pareto criterion, where different selves are treated as different individuals (see, e.g., Goldman 1979; Laibson 1997; Diamond and Köszegi 2003). This approach has the caveat of leading to an incomplete ranking of allocations. Most papers, however, have treated the agent's long-run preferences as the relevant for social welfare (e.g., Akerlof 1991; Gruber and Köszegi 2001; O'Donoghue and Rabin 2003; DellaVigna and Malmendier 2004; Amador, Werning, and Angeletos 2006). This approach is defended by O'Donoghue and Rabin (1999) and Bernheim and Rangel (2005).

Here, I will follow most of the literature in treating the agent's long-run time preferences as the relevant for social welfare. Moreover, since partially naive consumers have a mistaken perception about their true time-consistency parameter, I will attribute the correct parameter β for the preferences used in welfare comparisons.

An equilibrium will be called efficient if it is not possible to benefit both the *long-run self* of consumers and firms (with at least one of them being strictly better off). This is the standard notion of Pareto efficiency given the assumption that the long-run time preferences using the correct discount parameter β reflect the consumer's "true preferences."

PROPOSITION 1: *Suppose consumers are sophisticated. Then the equilibrium is efficient in the case of investment goods* ($b > 0$) *and inefficient in the case of leisure goods* ($b < 0$).

Proof: The first part follows from the fact that, in the case of investment goods, the equilibrium maximizes the consumer's (long-run) utility subject to firms not getting negative profits. The second part follows from the fact

that, by the strict concavity of total surplus, the unique maximum is given by Equation (1). Since profits are equal to zero and the consumer is worse off under marginal-cost pricing, it follows that the equilibrium is inefficient. ■

Because there is no *ex ante* asymmetric information, regulation can implement the allocation that maximizes total surplus (first-best).[12] Furthermore, there are many ways of achieving the first-best allocation. I will focus on regulation through a sales tax financed by lump-sum transfers from consumers both because these are the instruments often used in practice (e.g., cigarette taxes) and because its intuition is straightforward.

The inefficiency in the case of leisure goods stems from the fact that, with marginal-cost pricing, consumers do not internalize the full effects of consuming the good. Since the consumers internalize only a fraction β of future costs $-\delta b$, they do not account for the remaining part: $-\delta b(1-\beta)$. Therefore, an optimal sales tax should increase the price of the good by $-\delta b(1-\beta)$ so that perceived costs equal perceived benefits. If the proceedings of the tax are returned lump-sum to consumers, their (long-run) utility increases.[13] Furthermore, the firms' profit remains equal to zero. Hence, the resulting allocation dominates the equilibrium without taxes. The tax increase serves as a commitment device that makes the long-run self better off.

For investment goods, the equilibrium is inefficient when consumers are partially naive because they do not take into account the true parameter β.[14] Let τ denote the sales tax and define the after-tax cost of providing the good as $\hat{a} := a + \tau$. The optimal tax on investment goods should be such that after-tax prices are equal to the price that maximizes total surplus: $a - \delta b(1-\beta)$.[15] Hence,

$$p = \hat{a} - \delta b(1-\hat{\beta})\frac{f(\hat{\beta}\delta b - p)}{f(\beta\delta b - p)} - \frac{F(\hat{\beta}\delta b - p) - F(\beta\delta b - p)}{f(\beta\delta b - p)} = a - \delta b(1-\beta)$$

Expressing in terms of τ, we obtain the sales tax that maximizes total surplus:

$$\tau = \delta b\left[(1-\hat{\beta})\frac{f(\hat{\beta}\delta b - p)}{f(\beta\delta b - p)} - (1-\beta)\right] - \frac{F(\hat{\beta}\delta b - p) - F(\beta\delta b - p)}{f(\beta\delta b - p)}.$$

[12]Since firms always get zero profits in equilibrium, total surplus is equal to consumer surplus.

[13]The same result obtains if the tax proceedings were given to firms conditional on providing access. In this case, consumers would extract these transfers through the lump-sum price L and, because usage prices are equal to marginal cost plus the sales tax, they would have no incentive to accept a second-period contract.

[14]More precisely, for generic distributions F, the equilibrium in the case of investment goods is inefficient.

[15]The aggregate surplus is given by $\pi(p) + F(\beta\delta b - p)(p - a)$.

I formally state the results above in the following proposition:

PROPOSITION 2: *If consumers are partially naive, then the equilibrium is inefficient. An efficient allocation can be obtained by imposing a sales tax of:*

(i) $\tau = -\delta b(1-\beta) > 0$ *in the case of case of leisure goods; and*

(ii) $\tau = \delta b[(1-\hat{\beta})\frac{f(\hat{\beta}\delta b - p)}{f(\beta\delta b - p)} - (1-\beta)] - \frac{F(\hat{\beta}\delta b - p) - F(\beta\delta b - p)}{f(\beta\delta b - p)}$ *in the case of investment goods.*

The allocation implemented by the sales tax maximizes the consumer surplus subject to the firms obtaining nonnegative profits.

In the case of sophisticated consumers, we obtain the following result.

COROLLARY 1: *If consumers are sophisticated, the sales tax that maximizes consumer surplus subject to firms obtaining nonnegative profits is:*

(i) $\tau = -\delta b(1-\beta) > 0$ *in the case of case of leisure goods; and*

(ii) $\tau = 0$ *in the case of investment goods.*

Note that the optimal sales tax for leisure goods is always positive. It is a function of the long-run discount factor δ, the time-inconsistency parameter β, and the expected delayed costs b of consuming the good and does not depend on the distribution of the immediate benefits c from consuming or the perceived degree of time-consistency $\hat{\beta}$.[16]

The optimal sales tax for investment goods may be either positive or negative. Furthermore, it requires not only knowledge of the long-run discount factor δ, the expected delayed costs b of consuming the good, and the time-inconsistency parameter β but also the distribution of c and the perceived parameter of time-inconsistency $\hat{\beta}$. Hence, determining the optimal tax for these goods requires knowledge of several parameters which are hard to be estimated.

4. Conclusion

Although the DellaVigna and Malmendier model appears to be successful in explaining evidence for the health club, vacation time-sharing, and credit cards industries, it is not compatible with evidence from some of the most compelling examples of markets featuring time-inconsistent consumers: tobacco, alcohol, and unhealthy food. According to their model, firms should be offering ex ante lump-sum transfers to consumers and higher unit prices so that consumers purchase exclusively from them.

[16]Gruber and Köszegi (2001) provide estimates of all the parameters required to calculate the optimal tax in our model for the case of cigarettes.

I have shown that the evidence can be explained if one drops the exclusivity assumption, which seems to be unreasonable in these markets. While in the exclusive-contracts case the difference between leisure and investment goods is only a sign change in prices, there is an important asymmetry between them when contracts are nonexclusive.

The marginal-cost-pricing result for the case of leisure goods suggests that taxes should focus on leisure goods such as cigarettes, alcohol, and unhealthy food and not on investment goods, where the market may provide commitment devices efficiently (at least if consumers are sophisticated; otherwise, the optimal tax may be positive or negative and depends on parameters that are hard to estimate).

References

AMADOR, M., I. WERNING, and G.-M. ANGELETOS (2006) Commitment vs. flexibility, *Econometrica* **74**, 365–396.

AKERLOF, G. A. (1991) Procrastination and obedience, *American Economic Review* **81**, 1–19.

BECKER, G. S., and K. M. MURPHY (1988) A theory of rational addiction, *Journal of Political Economy* **96**, 675–700.

BERNHEIM, B. D., and A. RANGEL (2005) Behavioral public economics: Welfare and public policy analysis with non-standard decision makers, NBER Working Paper number 11518.

DELLAVIGNA, S., and U. MALMENDIER (2004) Contract design and self-control: Theory and evidence, *Quarterly Journal of Economics* **119**, 353–402.

DELLAVIGNA, S., and U. MALMENDIER (2006) Paying not to go to the gym, *American Economic Review* **96**, 694–719.

DEWATRIPONT, M. (1988) Commitment through renegotiation-proof contracts with third parties, *Review of Economic Studies* **55**, 377–390.

DIAMOND, P., and B. KÖSZEGI (2003) Quasi-hyperbolic discounting and retirement, *Journal of Public Economics* **87**, 1839–1872.

ELIAZ, K., and SPIEGLER, R. (2006) Contracting with diversely naïve agents, *Review of Economic Studies* **73**, 689–714.

ELLISON, G. (2006) Bounded rationality in industrial organization, Blundell, in *Advances in Economics and Econometrics: Theory and Applications*, R. Persson and W. K. Newey, eds. Cambridge: Cambridge University Press, pp. 143–174.

GABAIX, X., and D. LAIBSON (2006) Shrouded attributes, consumer myopia, and information suppression in competitive markets, *Quarterly Journal of Economics* **121**, 505–540.

GOLDMAN, S. (1979) Intertemporally inconsistent preferences and the rate of consumption, *Econometrica* **47**, 621–626.

GRUBER, J., and B. KÖSZEGI (2001) Is addiction 'rational'? Theory and evidence, *Quarterly Journal of Economics* **116**, 1261–1303.

HART, O., and J. TIROLE (1988) Contract renegotiation and coasian dynamics, *Review of Economic Studies* **55**, 509–540.

HEIDHUES, P., and B. KÖSZEGI (2005) The impact of consumer loss aversion on pricing, University of Bonn and Berkeley University.

KÖSZEGI, B. (2005) On the feasibility of market solutions to self-control problems, *Swedish Economic Policy Review* **12**, 71–94.

LAIBSON, D. I. (1997) Golden eggs, and hyperbolic discounting, *Quarterly Journal of Economics* **112**, 443–477.

O'DONOGHUE, T. D., and M. RABIN (1999) Doing it now or later, *American Economic Review* **89**, 103–124.

O'DONOGHUE, T. D., and M. RABIN (2003) Studying optimal paternalism, illustrated by a model of sin taxes, *American Economic Review* **93**, 186–191.

OSTER, S. M., and F. M. SCOTT MORTON (2005) Behavioral biases meet the market: The case of magazine subscription prices, *Advances in Economic Analysis & Policy* **5**, article 1, available at: http://www.bepress.com/bejeap/advances.

SHUI, H., and L. M. AUSUBEL (2005) Time inconsistency in the credit card market, University of Maryland.

SPIEGLER, R. (2006) Competition over agents with boundedly rational expectations, *Theoretical Economics* **1**, 207–231.