Forthcoming in *European Economic Review*

To be consumed with moderation *

Juan D. Carrillo

Columbia Business School
and CEPR

This version: December 2002

Abstract

We study the optimal behavior of a hyperbolic discounting agent who has incomplete information about his own preferences and can only learn them through consumption. We show that, even if moderate current consumption and moderate future consumption always dominates abstinence, the agent may optimally decide not to consume as a commitment device against inefficient learning that would lead to future excesses. This provides a rationale for why smokers, gamblers or compulsive buyers stick to second-best personal rules of behavior –such as "abstinence"– without invoking standard habit formation arguments. We also study how urges modify the strategy of the individual.

Keywords: Hyperbolic Discounting, Consumption, Abstinence, Self-Knowledge.

JEL Classification: A12, D83, D90.

---

*This paper is a short version of my working paper entitled "Self-control, moderate consumption, and craving". I thank M. Castanheira, M. Dewatripont, J. Miron, J. Tirole, P. Weil, two anonymous referees and participants at ULB's Behavioral Economics workshop, Milan and Berlin for comments. I am specially indebted to Isabelle Brocas for endless discussions. All errors are mine. Correspondence address: Juan D. Carrillo, Columbia Business School, Uris Hall - 812, 3022 Broadway, New York, NY 10027, USA, e-mail <jc2095@columbia.edu>.

1 Introduction

In modern behavioral economics, the observation that individuals overweigh the present –i.e. they exhibit a "salience/impatience" for current payoffs relative to distant ones– is identified as a major cause of self-control problems.[1] Under this type of preferences, the agent tends to overconsume goods that provide an immediate reward but affect negatively his future welfare, for example through a deterioration of health (drinking, smoking, overeating) or a decrease in wealth (compulsive credit card purchases, consumption of luxury goods, gambling).

In the type of situations previously mentioned, moderation seems to be less frequent than excess or abstinence. This claim, based on casual observation, constitutes the starting point of this research.[2] More precisely we develop a theory that relates the problem of self-control to the decision to forego valuable consumption opportunities.

There exist some simple theories that account for this observation. First, one can argue that, for these activities, even moderation is harmful. However, this is not convincing since many individuals who abstain openly state that moderate consumption, if sustainable, would be desirable. A second explanation often invoked is that radical abstinence serves as a focal point or rule of thumb, which is suboptimal but nevertheless more likely to be followed than the loosely defined concept of moderation.[3] This argument is intuitively very appealing, but it does not explain why it should be easier to commit to 'zero' rather than to a less suboptimal focal point like 'one every week' or 'three per day'. Third, one might simply state that moderation is not sustainable in the long run if the good is physically addictive. However, as stated by Herrnstein and Prelec (1992), an addictive good is characterized by the combination of two factors: negative internalities (i.e. a negative externality within a single person) and habit formation (i.e. a complementarity between past and present consumption). By contrast, the purpose of this paper is to explain abstinence relying exclusively on negative internalities. Why? Because we believe that negative internality is the only characteristic shared by all the activities previously mentioned.[4] Furthermore, because we also think that habit formation is sometimes over-emphasized: except for hard drugs (which are clearly not the object of this study), low

[1] Despite some experimental evidence gathered first in Psychology and now in Economics, there is still a fair amount of controversy among economists on whether individuals do really overweigh current payoffs (see e.g. Rubinstein, 2001). This paper has nothing to add to this debate and simply assumes this discounting.

[2] Note that we intentionally refer both to goods and activities that are subject to habit formation (tobacco, alcohol) and goods and activities that are not (compulsive buying, betting).

[3] For example, Ainslie (1992) argues that *"Alcoholics find that they cannot engage in 'controlled drinking' [... and] can only hope never to be lured across the bright line between some drinking and no drinking"* [p.169].

[4] One might argue that compulsive buying, gambling and overeating are also subject to habit formation, maybe not physical but of some other kind. The problem is that if we define as "habit forming" all the activities where we observe that agents are unable to moderate their behavior, then explaining abstinence becomes a tautology.

levels of consumption like for example 'two cigarettes a day' are unlikely to hook physically an individual.

The present work thus takes a different route. It shows that abstinence can be a rational decision even in situations where it is strictly dominated by moderate consumption from the viewpoint of the agent at any date. This conclusion rests on three building blocks. First, the agent has hyperbolic discounting preferences and cannot commit to the decisions that will be taken in the future. Second, he has imperfect knowledge about the pleasure derived by the consumption of the good. Third, he learns about his own preferences through consumption.

The intuition for our result is as follows. When the agent consumes, he discovers his preferences. Due to time-inconsistency, learning can be harmful. For instance, if the agent's taste for the good turns out to be sufficiently high, he will fall in a state of preferences where he wishes high consumption at the present date and moderation afterwards but, due to his inability to commit, ends up overconsuming in every period. So, if the expected welfare loss of future excesses caused by learning offsets the gain of optimal current consumption, the individual strictly prefers abstinence. In other words, ignorance is a commitment device against a future behavior inefficient from the current perspective. However, to avoid learning, the individual must adopt a second-best abstinence resolution.

In a second step, we introduce the possibility of urges, defined as stochastic positive shocks in the benefit of consumption. Not surprisingly, experiencing an urge increases the agent's willingness to consume. We show that an agent subject to urges but who is not currently experiencing one will be more willing to abstain than an agent who is never subject to urges. The idea is simply that the former is more concerned than the latter about his future misbehavior. As a result, he is more inclined to use strategic abstinence as a no-learning commitment device.

Several studies are related to ours. However, unlike our paper, they predict either moderate consumption or "happy abstinence". Carrillo and Mariotti (2000) were the first to suggest that a hyperbolic discounting individual may remain strategically ignorant when information can be inefficiently used by future incarnations of the agent.[5] In their paper, any amount of information can be costlessly acquired before the consumption decision. The model predicts that ignorance mitigates the self-control problem through a moderation in consumption.[6] Yet, our goal is precisely to explain why individuals either abstain or overconsume. Caillaud, Cohen

[5]See also Brocas and Carrillo (2000a,2001), Bénabou and Tirole (2001,2002) and the survey by Brocas and Carrillo (2000b).

[6]In fact, in their basic binary model the agent has the choice only between consuming and not, so 'moderation' cannot be distinguished from 'abstinence'. However, as argued in section 6.1, under a larger set of alternatives, an agent who stops learning will, on average, decrease his consumption but not necessarily abstain.

and Jullien (1996) analyze the behavior of an individual who can "self-restrain" his choices. More precisely the paper assumes that, for any given strategy, the agent can restrict the set of potential deviations to strategies that will be themselves followed (i.e. those for which he will not have a further incentive to deviate). Just as in Carrillo and Mariotti (2000), their resulting self-restricted equilibrium predicts avoidance of excesses through moderation. Last, Becker and Murphy (1988) study consumption of an addictive good in a standard exponential discounting model. Their paper shows the existence of multiple stable steady-state equilibria where agents either abstain or consume substantially. Naturally, their results can only be applied to goods subject to habit formation (i.e. goods for which the past level of consumption affects the marginal utility of current consumption). Besides, in their model, individuals who do not consume are happy abstainers whereas our goal is to explain abstinence for individuals who would prefer to consume moderately.

2 Consumption with time-inconsistent preferences

2.1 The basic model

We analyze the decision of a single agent (he) to consume a free good that procures an immediate pleasure but affects negatively his future welfare. More precisely, the agent lives three periods ($\tau \in \{1,2,3\}$) and can consume in the first two ($t \in \{1,2\}$). If, at date t, the agent decides to consume an amount x_t ($\geqslant 0$), then he obtains an *instantaneous* positive utility $\theta\, v(x_t)$, where θ (>0) represents his taste for consumption. However, consumption also entails a *one-period delayed* cost $-c(x_t)$. For analytical tractability, we restrict our attention to the CRRA class of utilities and to linear costs.

Assumption 1 *$c(x) = x$ and $\theta\, v(x) = \theta\, x^\gamma/\gamma$ with $\gamma \in (0,1)$.*

We depart from the neoclassical paradigm with the assumption that the agent has *dynamically inconsistent preferences*, in the sense of Strotz (1956). Basically, it implies that the period-to-period rate of discount falls monotonically or, stated differently, that the present is "overweighed" relative to the distant future.[7] We formalize this inconsistency in the simplest possible way. We assume that the individual attaches a weight 1 to the present period and a weight β (<1) to every future period.[8] Note that $\beta = 1$ is the standard case of time-consistent

[7] See Ainslie (1992) or Loewenstein and Prelec (1992) for a theoretical discussion and some empirical support.

[8] The literature usually employs a slightly more general formalization: period $t+\tau$ is discounted at a rate $\beta\delta^\tau$ from the perspective of date t (see e.g. the seminal modelling by Phelps and Pollak (1968) and, more recently, Akerlof (1991), Laibson (1997), O'Donoghue and Rabin (1999), Brocas and Carrillo (2000a, 2001) and Palacios-Huerta (2001) among others). Setting $\delta = 1$ implies no loss of generality. For a review of the different ways of modelling time-inconsistency we refer to Caillaud and Jullien (2000).

preferences. When $\beta < 1$, the agent plays a non-cooperative intrapersonal game. Naturally, as β decreases, the individual is less able to internalize the effects of his current decision on future welfare, so the intrapersonal conflict becomes more acute.

We denote by $V_t(\cdot)$ the intertemporal utility of the agent from the perspective of date t. Given the discount function and utilities described above, we have:

$$V_1(x_1, x_2; \theta) = \Big[\theta\, v(x_1) - \beta\, c(x_1)\Big] + \beta\Big[\theta\, v(x_2) - c(x_2)\Big] \quad (1)$$

$$V_2(x_1, x_2; \theta) = -c(x_1) + \Big[\theta\, v(x_2) - \beta\, c(x_2)\Big] \quad (2)$$

From the examination of (1) and (2) we can already obtain some preliminary insights. First, since we are not focusing on habit formation problems, the consumption decision at each period x_1 and x_2 is *independent of past and future behavior*. This allows us to isolate the effect of negative internalities on welfare given time-inconsistent preferences. Second, the dynamically inconsistency of preferences implies that the net payoff of consumption at date 2 from the perspective of date 2 ($\theta\, v(x_2) - \beta\, c(x_2)$ as given by (2)) does not coincide with the discounted net payoff of consumption at date 2 from the perspective of date 1 ($\theta\, v(x_2) - c(x_2)$ as given by (1)). As we show in the next section, this implies that optimal contingent plans at some date are no longer optimal when reconsidered one period later. Third, our modelling implies that *moderate current consumption and moderate future consumption strictly dominate abstinence*. To see this, note that, using Assumption 1, we have $\theta\, v(0) = \beta\, c(0)$ and $\theta\, v'(0) = \infty > \beta\, c'(0) = \beta$ (similarly, $\theta\, v(0) = c(0)$ and $\theta\, v'(0) = \infty > c'(0) = 1$). This puts ourselves in the worst possible scenario for observing a strategy of abstinence.

2.2 Consumption under full commitment and no commitment

The purpose of this work is to analyze the optimal behavior when the agent has incomplete information about his own preferences and consumption is informative. Imperfect self-knowledge is captured with the taste for consumption parameter θ, for which the individual only knows the distribution. However, in order to provide a benchmark for comparison, we first study the case where θ is fixed and known.

For any given θ, denote by $x^*(\theta)$ and $x^{**}(\theta)$ the optimal number of units that the individual desires to consume in the current period and at a future date, respectively. From (1), we have:

$$x^*(\theta) = \arg\max_x \ \theta\, v(x) - \beta\, c(x) \qquad \Rightarrow \qquad x^*(\theta) = \left(\frac{\theta}{\beta}\right)^{\frac{1}{1-\gamma}} \quad (3)$$

$$x^{**}(\theta) = \arg\max_x \ \theta\, v(x) - c(x) \qquad \Rightarrow \qquad x^{**}(\theta) = \theta^{\frac{1}{1-\gamma}} \quad (4)$$

Using (1), (2), (3) and (4), we can characterize the optimal and equilibrium consumption paths of the agent. If the individual could commit at date 1 on his present and future consumption behavior, he would choose $x_1 = x^*(\theta)$ and $x_2 = x^{**}(\theta)$. By contrast, if commitment on future behavior is not possible, the agent anticipates a consumption $x_2 = x^*(\theta)$ at date 2, which does not modify his optimal consumption $x_1 = x^*(\theta)$ at date 1. Obviously, the individual is always strictly better-off if he can decide his future consumption: $V_1(x^*(\theta), x^{**}(\theta); \theta) > V_1(x^*(\theta), x^*(\theta); \theta)$ for all θ.

Given $\beta < 1$, we have $x^*(\theta) > x^{**}(\theta)$ for all θ. That is, a dynamically inconsistent individual desires to consume more in the present period than in future periods. However, because of his inability to commit to a given consumption path, the agent ends up *overconsuming*. Also, note that the agent's willingness to consume is increasing in his taste θ: $\partial x^*/\partial\theta > 0$ and $\partial x^{**}/\partial\theta > 0$. The most interesting point is that the welfare of an individual who overconsumes may be strictly smaller than that of an agent who does not consume at all. This may occur even if, by assumption, the optimal level of current and future consumption is always above zero. In our model, the utility derived by a future (excessive) consumption is:

$$\beta\Big[\theta\, v(x^*(\theta)) - x^*(\theta)\Big] = \frac{\beta}{\gamma}\left(\frac{\theta}{\beta}\right)^{\frac{1}{1-\gamma}}(\beta - \gamma) \tag{5}$$

Therefore, abstinence dominates overconsumption if and only if:

$$\beta\Big[\theta\, v(x^*(\theta)) - x^*(\theta)\Big] < 0 \quad \Leftrightarrow \quad \beta < \gamma \tag{6}$$

that is, if the intrapersonal conflict is sufficiently important. Note also from (5) that, if $\beta < \gamma$, then the payoff will be decreasing in θ, the agent's taste for the good: a higher utility of consuming today implies a higher desire to overconsume tomorrow which results in a lower overall welfare. In fact, equation (5) captures the idea that repeated abstinence may be a second-best solution, but still better than systematic excesses. The problem is to understand how is it possible to sustain the no-consumption regime. Elster (1989) suggests that *"The trick is to put oneself in a frame of mind in which one violation of the rule allows one to predict rule violations on all later occasions. [...] By setting up this domino effect, I raise the stakes"* [p.49]. This idea has been formalized by Bénabou and Tirole (2001): a person with imperfect recall of past feelings can achieve self-control by the fear of transforming a lapse into a long run damaging precedent. In the next section, we point at a different mechanism that encourages abstinence.

2.3 Learning through consumption

We now analyze the consumption decision when the agent has imperfect knowledge about his own preferences, learns through consumption and cannot commit on his future choices. More specifically, we suppose that the agent's taste for consumption θ is drawn from a probability distribution with support $[\underline{\theta}, \overline{\theta}]$, density $f(\theta)$ and c.d.f. $F(\theta)$. The agent knows the distribution but ignores the realization of the parameter.[9]

Suppose first that consumption is uninformative. In that case, the individual behaves in both periods according to the expected value of his taste $E[\theta]$. More specifically, and following the argument developed in the previous section, the agent's consumption at dates 1 and 2 is:

$$x^*(E[\theta]) \;=\; \arg\max_x \;\int_{\underline{\theta}}^{\overline{\theta}} \Big[\theta\, v(x) - \beta\, c(x)\Big]\, dF(\theta) \quad \Rightarrow \quad x^*(E[\theta]) \;=\; \left(\frac{E[\theta]}{\beta}\right)^{\frac{1}{1-\gamma}} \qquad (7)$$

By contrast, when consumption allows the agent to learn about his own tastes, the optimal level of consumption is revised from period to period.

Learning is modelled in its most basic form. We suppose that if the agent consumes any strictly positive amount at date 1 ($x_1 > 0$), then he learns perfectly his taste θ for the good and uses this information for his consumption decision at date 2. If the agent decides not to consume at all ($x_1 = 0$), then no learning occurs.

This formalization is obviously unrealistic for several reasons. First, preferences may change over time. Second, it is unlikely that an individual will perfectly learn his tastes after only one period of consumption. Third, even if the agent abstains, he may receive some valuable information about the characteristics of the product. Last, and most importantly, learning should be proportional to the amount of consumption. Naturally, such a shortcut entails a loss of generality in our results. However, our goal is to show that if abstinence reduces "substantially" the amount of learning about the agent's own preferences (which seems a reasonable assumption), then there is scope for no-consumption: it induces ignorance which in turn reduces the likelihood of future overconsumption. The extreme modelling proposed in this paper provides the simplest way of capturing this trade-off between optimal consumption and inefficient learning.

Given this learning technology, the agent at date 1 has basically two alternatives. First, to consume the optimal amount from his current perspective and given his current knowledge $x_1 = x^*(E[\theta])$, anticipating that he will learn his exact taste and therefore that he will consume $x_2 = x^*(\theta)$ in the following period. Second, to refrain from consuming ($x_1 = 0$), in which case

[9] Imperfect self-knowledge is becoming standard in behavioral economics (see e.g. Brocas and Carrillo (2000a), Bénabou and Tirole (2002) or Bodner and Prelec (2002)).

no information is revealed so his consumption at date 2 will be $x_2 = x^*(E[\theta])$.[10] We denote by V_1^c and V_1^a the expected intertemporal welfare of the agent from his date-1 perspective if he consumes and if he abstains in the first period, respectively. Formally:

$$V_1^c = \int \Big[\theta\, v(x^*(E[\theta])) - \beta\, c(x^*(E[\theta]))\Big] + \beta\Big[\theta\, v(x^*(\theta)) - c(x^*(\theta))\Big] dF(\theta) \tag{8}$$

$$V_1^a = \int \beta\Big[\theta\, v(x^*(E[\theta])) - c(x^*(E[\theta]))\Big] dF(\theta) \tag{9}$$

The difference in intertemporal utility between optimal current consumption and abstinence is:

$$V_1^c - V_1^a = \int (1-\beta)\,\theta\, v(x^*(E[\theta])) + \beta\Big[\theta\, v(x^*(\theta)) - c(x^*(\theta))\Big] dF(\theta) \tag{10}$$

Equation (10) captures the fundamental trade-off of the paper. Current consumption has both benefits and costs. On the one hand, the individual enjoys the current utility of consumption instead of postponing it to a future date, as reflected by the first term in the r.h.s. of the equation. On the other hand, consumption triggers-off learning and opens the door to future overconsumption, as reflected by the second term in the r.h.s. of the equation. Naturally, the relative magnitude of these two effects depends on the inconsistency parameter. As β decreases, the individual is less able to internalize the effect of his future behavior on his welfare. He is therefore more likely to overconsume, which technically means that $\theta\, v(x^*(\theta)) - c(x^*(\theta))$ is more likely to be negative. At the same time, the individual puts a smaller weight on such inefficient future behavior. Last, note from (5) that the second term in (10) is negative if and only if $\beta < \gamma$. Therefore, a sufficient condition for consumption being more desirable than abstention (despite its future level being excessively high) is $\beta > \gamma$.

Using the expressions (5) and (7) we can rewrite (10) as follows:

$$V_1^c - V_1^a = \left(\beta^{\frac{-\gamma}{1-\gamma}}\gamma^{-1}\right)\left[(1-\beta)(E[\theta])^{\frac{1}{1-\gamma}} - (\gamma-\beta)E[\theta^{\frac{1}{1-\gamma}}]\right] \tag{11}$$

From (11) we deduce that the agent will strictly prefer first-period abstinence to first-period consumption if and only if:[11]

$$V_1^c - V_1^a < 0 \quad \Leftrightarrow \quad \beta < \bar{\beta} \equiv \frac{\gamma\, E[\theta^{\frac{1}{1-\gamma}}] - (E[\theta])^{\frac{1}{1-\gamma}}}{E[\theta^{\frac{1}{1-\gamma}}] - (E[\theta])^{\frac{1}{1-\gamma}}} \tag{12}$$

[10] The fact that learning is independent of x implies that consuming an amount positive but different from $x^*(E[\theta])$ is a strategy dominated by consuming $x^*(E[\theta])$. However, the result could be generalized to other learning technologies. Intuitively, if the amount of learning is sufficiently increasing in the quantity consumed, the two possible equilibria will be either to abstain or to consume an amount in the neighborhood of $x^*(E[\theta])$.

[11] Note that $1/(1-\gamma) > 1$ since $\gamma \in (0,1)$. The function $\theta^{\frac{1}{1-\gamma}}$ is then increasing and convex, which implies that $E[\theta^{\frac{1}{1-\gamma}}] > (E[\theta])^{\frac{1}{1-\gamma}}$.

The major conclusion we can draw from (12) is that an agent may rationally choose abstinence in order to avoid future excesses. For this to happen we need the combination of two factors. First, the intrapersonal conflict must be sufficiently important (β low). Recall from (5) that abstinence may dominate overconsumption. However, in order to be desirable, abstinence must also compensate for the foregone utility of current consumption. Therefore, the condition for the optimality of first-period abstinence (12) is more stringent than the condition in (6) or, stated formally, $\bar{\beta} < \gamma$. Second, the amount of uncertainty resolution between consumption and abstinence must be sufficiently high.[12] Other things being equal, as the spread of the taste parameter increases, the amount of learning through consumption increases and so does the benefit of ignorance as a device to avoid excesses (technically, $E[\theta^{\frac{1}{1-\gamma}}] - (E[\theta])^{\frac{1}{1-\gamma}}$ increases).

To sum up, we have provided a rationale for abstinence in situations where positive but moderate consumption would always be a superior strategy. Instead of imposing the "zero-resolution" as an ad-hoc focal point, we have built on the idea that abstinence substantially decreases how much the individual learns about his own preferences. Since, under time-inconsistent preferences, learning can induce a future level of consumption excessively high from the current perspective, abstinence then provides a commitment device against such inefficient behavior.

At this point, some further remarks are in order. First, increasing the number of consumption periods can only strengthen our conclusions. If the horizon is large but finite, then a standard backward induction argument proves that the agent will abstain in every period except the last one, as long as (12) holds. More interestingly, if the horizon is infinite (or stochastic) abstinence can be sustained more easily. Indeed, the individual is concerned with the possibility that present consumption might be followed by perpetual overconsumption. So, as long as the agent anticipates future abstinence, he is also willing to abstain currently. This suggests that our model may explain not only *convergence to* but also *persistence of* abstinence.[13] Second, including some learning under abstention would not modify our result. As explained above, our result relies on the trade-off between net utility of current consumption and potential inefficiency of future consumption caused by learning. As a result, insofar as consumption substantially increases the amount of information revealed, the decision to abstain can be sustained. Third, the model predicts that an individual who decides to abstain will minimize the chances of being exposed to information about his preferences. In particular, he will avoid both the people and the environments where this consumption is likely to play an important role.

[12] For instance, we know that first-period abstinence never occurs under perfect information or under no learning.

[13] For a formal analysis of the problem in an infinite horizon setting with a gradual resolution of uncertainty, see the longer version of this paper (Carrillo, 2001). There, we show that strategic abstinence as a no-learning commitment device holds quite generally. Moreover, we also prove other results such as the existence of multiple Pareto rankable equilibria, each of them characterized by a different amount of learning and consumption.

3 Urges

Individuals commonly state that, on specific occasions, an irresistible "force" induces them to deviate from the no-consumption strategy. Those stimuli may be divided into internal impulses (craving) or external factors (an invitation to a party). The first type of stimuli is more problematic as it entails a greater difficulty in evaluating objectively whether an urge was really present. However, even under objective assessment of urges, casual observation suggests that succumbing to a temptation may be highly damaging from a long run perspective.

The purpose of this section is not to provide a comprehensive treatment of consumption under urges, but rather to explore within our framework one particular effect of urges on the behavior of agents. We model urges in the simplest possible way: at each period and with exogenous probability p the agent receives a positive shock on the benefit derived by current consumption. This is formalized as follows.

Assumption 2 *The instantaneous utility under an urge is:* $\theta\, v_u(x) = a\,\theta\, x^\gamma/\gamma$ *with* $a > 1$.

We also rule out any intrapersonal problem of information transmission and assume that the individual knows perfectly before his consumption choice whether he is acting under an urge or not. Denote by $x_u^*(\theta)$ and $x_u^*(E[\theta])$ the analogue of expressions (3) and (7) when the individual is subject to an urge. Urges increase the optimal quantity consumed:

$$x_u^*(\theta) \;=\; \left(\frac{a\,\theta}{\beta}\right)^{\frac{1}{1-\gamma}} \qquad \text{and} \qquad x_u^*(E[\theta]) \;=\; \left(\frac{a\,E[\theta]}{\beta}\right)^{\frac{1}{1-\gamma}} \tag{13}$$

Following the same methodology as in section 2, we denote by $V_{1;u}^c$ and $V_{1;n}^c$ the intertemporal utility of consumption when the agent currently experiences and does not experience an urge:

$$\begin{aligned} V_{1;u}^c \;=\; & \int \Big[\theta\, v_u(x_u^*(E[\theta])) - \beta\, c(x_u^*(E[\theta]))\Big] + \beta\, p\Big[\theta\, v_u(x_u^*(\theta)) - c(x_u^*(\theta))\Big] \\ & + \beta\,(1-p)\Big[\theta\, v(x^*(\theta)) - c(x^*(\theta))\Big] dF(\theta) \end{aligned} \tag{14}$$

$$\begin{aligned} V_{1;n}^c \;=\; & \int \Big[\theta\, v(x^*(E[\theta])) - \beta\, c(x^*(E[\theta]))\Big] + \beta\, p\Big[\theta\, v_u(x_u^*(\theta)) - c(x_u^*(\theta))\Big] \\ & + \beta\,(1-p)\Big[\theta\, v(x^*(\theta)) - c(x^*(\theta))\Big] dF(\theta) \end{aligned} \tag{15}$$

Since future urges are independent of current behavior, in our model the intertemporal utility of the agent under abstinence is the same whether he currently experiences an urge or not:[14]

$$V_{1;u}^a = V_{1;n}^a = \int \beta p\Big[\theta\, v_u(x_u^*(E[\theta])) - c(x_u^*(E[\theta]))\Big] + \beta(1-p)\Big[\theta\, v(x^*(E[\theta])) - c(x^*(E[\theta]))\Big] dF(\theta) \tag{16}$$

[14] One might argue that abstinence is more difficult and therefore yields lower utility under an urge. Also, the mere fact of being subject to a current urge conveys some information about one's preferences (formally, p should not be exogenous). Our stylized model does not capture these interesting properties.

Using (13), (14), (15) and (16), we can determine the conditions under which first-period abstinence is the optimal strategy of the individual.

$$V^c_{1;u} - V^a_{1;u} \;<\; 0 \quad \Leftrightarrow \quad \beta \;<\; \bar{\beta}_u \;\equiv\; \frac{\gamma\Big[E[\theta^{\frac{1}{1-\gamma}}] - y\Big] - (E[\theta])^{\frac{1}{1-\gamma}}}{\Big[E[\theta^{\frac{1}{1-\gamma}}] - y\Big] - (E[\theta])^{\frac{1}{1-\gamma}}} \tag{17}$$

$$V^c_{1;n} - V^a_{1;n} \;<\; 0 \quad \Leftrightarrow \quad \beta \;<\; \bar{\beta}_n \;\equiv\; \frac{\gamma\Big[E[\theta^{\frac{1}{1-\gamma}}] + z\Big] - (E[\theta])^{\frac{1}{1-\gamma}}}{\Big[E[\theta^{\frac{1}{1-\gamma}}] + z\Big] - (E[\theta])^{\frac{1}{1-\gamma}}} \tag{18}$$

where $y = (1-p)\frac{a^{\frac{1}{1-\gamma}}-1}{a^{\frac{1}{1-\gamma}}}\Big(E[\theta^{\frac{1}{1-\gamma}}] - (E[\theta])^{\frac{1}{1-\gamma}}\Big)$ and $z = p(a^{\frac{1}{1-\gamma}} - 1)\Big(E[\theta^{\frac{1}{1-\gamma}}] - (E[\theta])^{\frac{1}{1-\gamma}}\Big)$. Although these expressions are quite messy, we can obtain from them an interesting insight:

$$\bar{\beta}_u \;<\; \bar{\beta} \;<\; \bar{\beta}_n \quad (<\; \gamma) \tag{19}$$

Urges modify the consumption strategy of the individual in two ways. Their existence increases the likelihood and damage of a future misbehavior. Therefore, conditional on not being currently under an urge, the anticipation of a possible future urge increases the willingness of the forward-looking agent to refrain from consuming, as a commitment against future excesses ($\bar{\beta}_n > \bar{\beta}$). However, if the agent is currently under an urge, the desire of current satiation will offset the incentives to avoid learning, and consumption will be more likely to take place ($\bar{\beta}_u < \bar{\beta}$).

Overall, recall that in our model there is no habit formation. Therefore, under no learning, the individual does not condition his current decision on the likelihood and effect of future desires and future behavior (formally, if there is no current urge, $x_1 = x^*(E[\theta])$ for all p and a). The possibility to discover the taste for consumption modifies this conclusion: the agent is more reluctant to consume under no urge because he anticipates that excesses under future urges can be extremely self-damaging.

We have focused on the analysis of the agent's optimal behavior. However, one might wonder whether the very presence of urges is beneficial or detrimental for welfare. To analyze this issue denote by $V^c_{1;e}$ the expected intertemporal benefit of consumption before the agent knows whether there is a current urge or not. That is, $V^c_{1;e} = p\,V^c_{1;u} + (1-p)V^c_{1;n}$. Also, given $V^a_{1;u} = V^a_{1;n}$, the expected intertemporal benefit under abstinence is $V^a_{1;e} = V^a_{1;u} = V^a_{1;n}$. Using (13), (14), (15) and (16) we have:

$$\frac{\partial V^c_{1;e}}{\partial a} \;\gtrless 0 \quad \Leftrightarrow \quad \beta \;\gtrless\; \gamma - (1-\gamma)\frac{(E[\theta])^{\frac{1}{1-\gamma}}}{E[\theta^{\frac{1}{1-\gamma}}]} \qquad \text{and} \qquad \frac{\partial V^a_{1;e}}{\partial a} \;\gtrless 0 \quad \Leftrightarrow \quad \beta \;\gtrless\; \gamma \tag{20}$$

The implications are rather straightforward. When the intrapersonal conflict is sufficiently small ($\beta > \gamma$), the individual finds it optimal to consume in both periods. Moreover, despite the fact

that future consumption is excessive from the current perspective, it is still desirable. Hence, a higher benefit of consumption (i.e. an increase in a) translates into a higher welfare for the agent. On the opposite extreme, when the intrapersonal conflict is very acute ($\beta < \bar{\beta}_u$), the agent optimally abstains. Due to his future excessive consumption, the individual's welfare will then be lower the higher the urge a. For intermediate values of β, the agent's consumption decision depends on whether there is an urge or not. The two effects previously described operate simultaneously, so welfare may increase or decrease with the intensity of the urge depending on the parameters of the model. Naturally, a similar reasoning holds when we consider an increase in p, the likelihood of urges.

Remark 1. Note that we have adopted the most favorable modelling of urges concerning its effect on the well-being of the individual: an urge increases the benefit of current and future consumption without affecting the cost ($\partial v_u/\partial a > 0$ and $\partial c/\partial a = 0$) whereas the utility under no urge is the same as before. Even in that case, we have shown that the existence of urges can be welfare damaging. Naturally, with a more neutral modelling of urges (like, for example, a utility equal to $a\,\theta\,x^\gamma/\gamma$ under urge and $b\,\theta\,x^\gamma/\gamma$ under no urge with $p\,a + (1-p)b = 1$) this conclusion would only be strengthened.

Remark 2. Our model has emphasized abstinence as a mechanism to avoid learning. However, one can think of situations in which the trade-off is reversed, i.e. where the agent inefficiently consumes a small amount of the good only to learn about his own preferences. For example, suppose that there are two cutoffs θ_l and θ_h (with $\theta_l < \theta_h$) such that under an urge the agent consumes if $\theta > \theta_l$ and under no urge the agent consumes if $\theta > \theta_h$, otherwise he happily abstains.[15] Now, assume that $E[\theta]$ is slightly above θ_l but below θ_h and there is currently no urge. The agent may choose to consume some positive amount in the hope of learning some good news (θ smaller than θ_l) and therefore avoid harmful overconsumption in the event of experiencing an urge in the future.

4 Conclusion

The purpose of this paper has been to explore why individuals do not moderate the consumption of goods subject to self-control problems. We have studied the strategy of a dynamically inconsistent agent who has incomplete information about his own preferences and can only learn them through consumption. We have shown that, even if moderate current consumption and moderate future consumption is always better than abstinence, the agent may decide not to

[15]This would require a model different from the one presented in the paper where, by assumption, some consumption is always optimal.

consume as a commitment against inefficient learning that would lead to future excesses. We have also analyzed how urges affect the behavior and welfare of the agent. In particular, we have shown that anticipating the existence of future urges makes the individual more willing to abstain, provided that he faces no immediate necessity.

We would like to conclude with some suggestions for future research. First, urges are modelled in an extremely stylized way (see e.g. footnote 14). In our view, the type of urges that are more pervasive and, at the same time, more difficult to deal with are internal impulses. It would be interesting to study how and why craving acts in the motivational state of the individual so that he subjectively reevaluates the costs and benefits of consumption. Second, the value of ignorance has extensively been recognized in economics. It is therefore not surprising that an individual with internally conflicting interests may also find desirable to keep an imperfect self-knowledge. Issues related to self-esteem, or more generally to self-perception, could and should be analyzed under this light. Last, it should be possible to determine empirically which goods are more likely to be subject to habit formation. This in turn would provide compelling evidence on whether time-inconsistency is a determinant factor in the agents' decision to abstain.

References

1. Ainslie, G., 1992, Picoeconomics (Cambridge University Press).
2. Akerlof, G., 1991, Procrastination and Obedience, American Economic Review 81, 1-19.
3. Becker, G.S. and M.K. Murphy, 1988, A Theory of Rational Addiction, Journal of Political Economy 96, 675-700.
4. Bénabou, R. and J. Tirole, 2001, Willpower and Personal Rules, mimeo.
5. Bénabou, R. and J. Tirole, 2002, Self-Confidence and Personal Motivation, Quarterly Journal of Economics 117(3), 871-915.
6. Bodner, R. and D. Prelec, 2002, The Diagnostic Value of Actions in a Self-Signaling Model, in: I. Brocas and J.D. Carrillo, eds., The Psychology of Economic Decisions, Vol.1 (Oxford University Press) (in press).
7. Brocas, I. and J.D. Carrillo, 2000a, Entrepreneurial Boldness and Excessive investment, CEPR D.P. 2213.

8. Brocas, I. and J.D. Carrillo, 2000b, The Value of Information when Preferences are Dynamically Inconsistent, European Economic Review 44, 1104-1115.

9. Brocas, I. and J.D. Carrillo, 2001, Rush and Procrastination under Hyperbolic Discounting and Interdependent Activities, Journal of Risk and Uncertainty 22, 141-164.

10. Caillaud, B. and B. Jullien, 2000, Modelling Time-Inconsistent Preferences, European Economic Review 44, 1116-1124.

11. Caillaud, B., D. Cohen and B. Jullien, 1996, Towards a Theory of Self-Restraint, mimeo.

12. Carrillo, J.D., 2001, Self-control, Moderate Consumption, and Craving, CEPR D.P. 2017.

13. Carrillo, J.D., and T. Mariotti, 2000, Strategic Ignorance as a Self-Disciplining Device, Review of Economic Studies 67, 529-544.

14. Elster, J., 1989, Nuts and Bolts for the Social Sciences (Cambridge University Press).

15. Herrnstein, R. and D. Prelec, 1992, A Theory of Addiction, in: G. Loewenstein and J. Elster, eds., Choice over Time (Russell Sage Foundation, New York).

16. Laibson, D.I., 1997, Golden Eggs and Hyperbolic Discounting, Quarterly Journal of Economics 112, 443-477.

17. Loewenstein, G. and D. Prelec, 1992, Anomalies in Intertemporal Choice: Evidence and an Interpretation, Quarterly Journal of Economics 57, 573-598.

18. O'Donoghue, T. and M. Rabin, 1999, Doing it Now or Later, American Economic Review 89, 103-124.

19. Palacios-Huerta, I., 2001, Hyperbolic Discounting, Habits and Psychological Expected Utility Theory, mimeo.

20. Phelps, E.S. and R.A. Pollak, 1968, On Second Best National Saving and Game-Theoretic Growth, Review of Economic Studies 35, 185-199.

21. Rubinstein, A., 2001, Is it 'Economics and Psychology'? The case of Hyperbolic Discounting, mimeo.

22. Strotz, R.H., 1956, Myopia and Inconsistency in Dynamic Utility Maximisation, Review of Economic Studies 23, 166-180.