ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

SCHOOL OF ECONOMIC SCIENCES

DEPARTMENT OF ECONOMICS

DEPARTMENT OF INTERNATIONAL & EUROPEAN ECONOMIC STUDIES

Approaches to Bounded Rationality: The case of Time Inconsistency

Aris Valasiou

Dissertation submitted

in partial fulfilment of the necessary prerequisites

for the acquisition of the MSc Degree

Athens

March 2020

Supervisor: Nikolaos Vettas, Athens University of Economics and Business

Examiner 1: Christos Genakos, Athens University of Economics and Business

Examiner 2: Theodore Palivos, Athens University of Economics and Business

Table of Contents

Abstract

Perfect rationality is the cornerstone upon which the foundation of economics as a science is based. However, there is an ongoing trend among the economic theory literature to relax this assumption, based on both experimental and empirical evidence. In this dissertation, the focus is placed on a specific case of bounded rationality models involving time-inconsistent agents, who put disproportional weight to immediate outcomes relative to delayed ones. We overview the experimental evidence suggesting that agents are time-inconsistent and have self-control problems, as well as the most popular ways to model them, such as the (β,δ) model. An important aspect of those models is the distinction between 'sophistication' and 'naïvete', regarding individuals' cognizance of their self-control problems. Also, a plethora of observations into which time-inconsistency helps us to gain useful insights is presented, such as the tendency of individuals to procrastinate, or their mispredictions about future health club attendance. Special attention is given to the case of addiction, where a fully rational and a boundedly rational model are presented vis-à-vis, and their policy implications are discussed. A significantly higher excise tax on cigarette consumption is justified if agents are assumed to be time-inconsistent, as well as other measures that work as commitment devices for them. Lastly, we briefly discuss the theoretical motivation behind the bounded rationality hypothesis, and address some critiques aimed at it, such as those that view as appropriate the use of 'rationalized' models to describe human behavior.

KEYWORDS: Bounded Rationality, Time-Inconsistency, Self-Control, Naïvete

Chapter 1

Introduction

Perhaps the most important assumption that economic science is based upon is the assumption of rationality. Rational agents have three main characteristics, as Rubinstein (1991) notes: *i)* they know what is feasible, i.e. what elements constitute their choice set; *ii)* they know what is preferable, i.e. they have a complete preference relationship over their choice set; and *iii)* they know how to achieve the optimal (according to their preferences) outcome over the feasible ones. Of course, the notion of rationality was the cornerstone of nearly every major achievement that economic science has made, and its rapid expansion attests to the importance of the rationality paradigm, despite the fact that it seems somewhat far-fetched. A multitude of studies has shown that individuals in many settings act in ways contradicting rational behavior. Even Gary Becker, one of the most fervent advocates of the rationality paradigm, conceded in his Nobel speech (1993) that "actions are constrained by income, time, *imperfect memory, calculating capacities* and other limited resources".

The first reference to non-rational behavior is due to Simon (1955, 1956), who develops the notion of 'satisficing'. Simon hypothesizes that boundedly rational agents, instead of exhaustingly search all available options in order to make the best decision possible, perform a limited search until they can find a somewhat satisfactory decision, or one that suffices to satisfy their needs, hence the term satisficing. However, he does not provide us with an explicit criterion on which this procedure is based in order to model it appropriately.

Kahneman & Tversky (1979), in one of the most cited articles in the whole economics discipline, were the first to propose a model motivated by non-rational behavior. In an attempt to explain some puzzling findings on choices involving risk, such as the Allais paradox (see Allais (1953)), that seemed to contradict expected utility theory, the extension of rational decision making under uncertainty, they developed what they call 'prospect theory'. Under prospect theory, the choice between lotteries or prospects is characterized by *i)* reference dependence (all outcomes are evaluated against a reference point, not overall wealth); *ii)* loss aversion (losses against that reference point are averted compared to equivalent gains, hence symmetric fair bets are unattractive); *iii)* diminishing sensitivity (the marginal value of both gains and losses decreases as we get further from the reference point); and probability weighting (outcomes with a small probability to occur are overweighted, and those with a large one are underweighted).

Many articles reviewing the relevant literature exist (e.g. DellaVigna (2009), or Ellison (2006) in the context of Industrial Organization). We could group bounded rationality models in three categories, according to which one of the three main characteristics of the rationality paradigm is violated: models where agents cannot fully evaluate their choice set, models where agents have non-standard preferences, and models where agents engage in non-optimal decision-making.

In the first category, we have models involving agents that face limitations on their ability to examine all the available choices. Examples include models of sample-based reasoning, where agents evaluate only realizations of random variables, without taking into account their true probability of happening, and extrapolate naïvely based upon them[1] (e.g. Spiegler (2006)), and models of coarse reasoning, where agents create a 'coarse representation' of the available data, grouping non-similar situations into categories, and make inferences treating them as equivalent (see e.g. Mullainathan et al. (2008)).

In the second category, we have models involving agents whose preferences differ from those usually assumed for rational agents, for example discounted utility theory for intertemporal choices or expected utility theory for choice under risk. Models falling in the second category include those who involve agents with time-inconsistent preferences, who put greater weight to immediate outcomes compared to delayed ones (e.g. Laibson (1997), Gul & Pesendorfer (2001)), or those with reference point effects, where outcomes are evaluated against a reference point (e.g. Kahneman & Tversky (1979), as discussed already, or Kőszegi & Rabin (2006)). Examples of reference point effects include the endowment effect (the observation that individuals evaluate differently an object depending on whether they are the seller or the buyer, as indicated by the discrepancy in WTP (willingness to pay) and WTA (willingness to accept), which was illustrated e.g. by Kahneman, Knetsch and Thaler (1990) in their classic mug experiment), or the status quo bias (the observation that individuals tend

[1] For instance, in many cases people use anecdotal evidence and overweight the probabilities of events such as personal experiences or stories that they heard.

to adhere to the status quo more frequently compared to what the standard rationality theory predicts, which was illustrated e.g. by Samuelson & Zeckhauser (1988)).

Finally, models falling in the third category involve agents who have limited ability to choose the optimal (according to their preferences) among the available options, resorting to simplified procedures called ‘heuristics’ or rules-of-thumb, for their decision-making. Examples of such models include framing models, where the assessment of an option depends on whether it is framed narrowly (irrelevant of the other ones available) or broadly (in contrast with the other options) (as in the well-known experiment in Kahneman & Tversky (1981), for a relevant model see e.g. Read, Loewenstein & Rabin (1999)) or because the agent ranks them along some of their dimensions rather than assessing them holistically (e.g. Eliaz & Spiegler (2011)). Satisficing, as conceived by Simon (1955), also falls within that category of models.

Obviously, the extended literature with respect to bounded rationality cannot be covered appropriately in the context of this work. We will limit our attention instead to the case of time-inconsistency, both providing evidence that support the notion that individuals do not treat intertemporal choices the way it would be expected by a rational man (for example discounting future outcomes at a constant rate), thereby violating key tenets of rational behavior, and presenting the most popular ways through which time-inconsistency is modeled by the relevant literature. In addition, we will briefly discuss the policy implications that follow from assuming time-inconsistent agents in the context of addictive goods, and whether generally bounded rationality modeling is useful, addressing the ‘can we get the same with a standard model’ critique.

We continue as follows: In chapter 2, we provide laboratory evidence that suggest that individuals are time-inconsistent. In chapter 3, we discuss popular models incorporating time-inconsistent agents, and check how many empirical findings, both in market and non-market settings, can be explained by time-inconsistency. In chapter 4, we present both a rational and a boundedly rational model in the case of addiction, in order to compare the different policy prescriptions that emerge. This will motivate us to present a more theoretical approach on bounded rationality and its usefulness in the last chapter, which concludes.

Chapter 2

Experimental evidence regarding intertemporal choice

2.1 The discounted utility (DU) model

The subject of intertemporal choice is a matter that has troubled economists for many years, starting from Rae (1834), who tried to attribute the difference in the wealth among nations to the 'desire of accumulation', a psychological factor determining a nation's proclivity towards savings and investment. Many more economists followed[2], since many aspects of human life involve decisions where the benefits are not reaped at the same time that the relevant costs are incurred. Samuelson (1937) proposed the first and, even to this day, most-widely used model to specify agents' intertemporal choices, the discounted utility (DU) model. According to the model, an agent's intertemporal utility function is specified as: $U_t = \sum_{i=0}^{\infty}(\frac{1}{1+\rho})^i \, u(c_{t+i})$, where $u(c_t)$ is the agent's instantaneous utility, derived from his consumption level c_t, and ρ is his rate of time preference. The parameter ρ is intended to capture the willingness of the agent to substitute current for future consumption.

Samuelson (1937) recognized that the DU model could be subjected to valid critiques. He admitted that the model holds neither descriptive ('It is completely arbitrary to assume that the individual behaves so as to maximize an integral of the form envisaged in [the DU model]') nor normative ('any connection between utility as discussed here and any welfare concept is disavowed') validity. Nevertheless, the model was quickly adopted as the basic framework for analyzing intertemporal choices, mainly due to its simplicity and facilitation of analytical derivations.

As Frederick, Loewenstein and O'Donoghue (2002) point out, this particular functional form makes some questionable implicit assumptions regarding agents' behavior. First of all, it is implied that two consumption sequences are valued only by the global utility they provide, i.e. the discounted sum of each period's consumption utility, but not by the distribution of the 'instantaneous' utilities across time. Hence,

[2] See e.g. Bohm-Bawerk (1889)

with a sufficiently high discount factor, a consumption profile with positive consumption in only one time period is treated as equivalent with one that yields the same global utility, but prescribes constant consumption over time. That way the possibility that agents have preferences for utility patterns, beyond the ones for consumption levels, is barred .[3]

Another implicit assumption is the one of consumption independence. According to that assumption, preferences over consumption profiles do not depend on the common parts between them, i.e. the utility experienced in one period by a consumption choice is not affected by the consumption choices made in the previous or next period. Consider for example two consumption sequences, where the choices are either stake or salad for meal, which are identical in all time periods except for t, with the first sequence prescribing steak and the second prescribing a salad for meal at time t. It is somewhat far-fetched to assume that these sequences are ranked only by the agent's instantaneous preference over these two choices, and whether the last meal of hers was salad is an irrelevant consideration. As Koopmans (1960) pointed out: '[…] there is no clear reason why complementarity of goods could not extend over more than one time period'.

Finally, the stronger assumption behind the DU model is that of constant discounting. Constant discounting implies that the discount rate between any two consecutive periods is constant. To be more precise, a discount function, $D(t), \forall\, t,$ can be used, describing the discount rate between periods t and t+1. In the case of the DU model, $D(t) = \frac{1}{1+\rho}, \forall\, t$, so the discount rate is constant and equal between any two consecutive periods. Thus, the only factor determining the intertemporal welfare derived from two sets of outcomes providing the same instantaneous utility is the distance between the times they are actualized, not the initial time of actualization, so as the delay of both outcomes for a common time interval should not change their ranking. In that case, a person's intertemporal preferences can be summarized with a single discount rate, and constant discounting implies time consistency. A person's intertemporal preferences are termed time-consistent if, for any two consumption

[3] For instance, a preference for a flat utility profile instead of an oscillatory one, when both yield the same global utility.

profiles $C_t = (c_t, \dots, c_T)$ and $C'_t = (c'_t, \dots, c'_T)$, with $c_t = c'_t$,

$$U_t(C_t) \geq U_t(C'_t) \quad iff \quad U_{t+1}(C_{t+1}) \geq U_{t+1}(C'_{t+1}).$$

2.2 Against the DU model: Hyperbolic discounting

A considerable amount of laboratory evidence exist stressing out the ways the DU model fails to provide a realistic descriptive account of human preference. The most easily assailable assumption of the DU model based on the existing empirical evidence is the one of constant discounting. Rather, hyperbolic discounting seems to be a more well-justified assumption. Hyperbolic discounting occurs when an agents' intertemporal preferences exhibit a declining discount rate, or more formally that the discount function $D(t)$ is decreasing in t.

In order to examine whether discount rates decline over time, many studies have been conducted asking the subjects what amount of money, expected to be received at a future time, would make them indifferent to receiving a certain dollar amount today. By calculating the implied discount rates, one would expect them to fall as time intervals grew larger if hyperbolic discounting were to be true. For instance, in a highly cited study, Thaler (1981) divided college students into three subsets, and asked them what amount of money they required in order to postpone the receipt of a prize, so as to be indifferent to receiving that prize today. Prizes and delays were different for each of the three subsets of participants (who were asked to fill a 3x3 table, comparing three monetary outcomes with three receipt delay intervals) and the results are presented in Table 1 (each of the first three forms refers to one subset). Thaler finds, for a $15 prize for instance, the implied average annual discount rates to be 345 percent over one-month, 120 percent over one-year and 19 percent over ten-year horizons, based on the median responses of $20, $50 and $100 for a deferral of one month, one year and ten years respectively[4], and a similar pattern is observed across the first three forms.

In many cases, agents exhibiting hyperbolic discounting are time-inconsistent, leading to the observed pattern of 'time preference reversal' found in many empirical

[4] These discount rates are calculated assuming continuously compounded discounting. So the solutions of the equations $15=$20*exp(-ρ/12), $15=$50*exp(-ρ) and $15=$100*exp(-10*ρ) are the implied one-month, one-year and ten-year annual discount rates, respectively.

studies. In most studies, subjects are presented with a choice between dated prizes (x, t), i.e. ordered pairs with a dollar amount and the receipt date for that amount as their first and second elements, respectively. Many studies have indicated that, when facing two choices of the form: $[(x_p, t_p)$ or $(x_s, t_s)]$ and $[(x_p, t_p + d)$ or $(x_s, t_s + d)]$[5], with $x_p < x_s$, and d a common delay in both payments, subjects may prefer the first option (the SS option, standing for Smaller reward, Sooner received) in the first choice set, but the second option (the LL option, standing for Larger reward, Later received) in the second one, thereby 'reversing' their preferences[6]. In essence, the subjects' willingness to delay a payment in return for interest is diminished the more immediate it becomes. For example, Ainslie & Haslam (1992) find that while a majority of subjects prefers a \$100 check today compared to a \$200 check to be received in 2 years, a majority also prefers the \$200 check cashed in 8 years compared to the \$100 one cashed in 6 years. The aforementioned examples are indicative of a multitude of studies finding similar results, thus providing support for hyperbolic discounting and time-inconsistency.

It should be noted however that certain studies have cast doubt in the hyperbolic discounting hypothesis. Two such examples come from Read (2001) and Rubinstein (2003). On the one hand, Read (2001) suggests that 'subadditive discounting' can also explain the most important evidence in favor of hyperbolic discounting, namely the declining implied discount rates. Subadditive discounting describes the increase in discount rates as the relevant time interval is partitioned in smaller subintervals. For a more formal illustration, Read defines $f_{T,T'} = \frac{V_T}{V_T'}$ to be the discount fraction between times T and T' $(T < T')$, where V_T and V_T' are the dollar amounts in times T and T' that make an individual indifferent between them. This discount fraction serves as a discount factor for any interval (the inverse of the discount rate ρ, according to the previous notation). Also, $f_{0,T,T'} := f_{0,T} \times f_{T,T'}$. Then, for example, subadditive discounting holds if $f_{0,T} > f_{0,T/2,T} = f_{0,T/2} \times f_{T/2,T}$ (1), while hyperbolic discounting

[5] P standing for present date, S for subsequent date

[6] Although subjects are asked to compare these options at the same time, it is implied that those who experienced this preference reversal would instead prefer the first option in the first choice set, if they were brought back to the lab d days later. As long as subjects know about this reversal and want to avoid it, they might seek commitment devices, and eagerness for commitment could serve as evidence in favor of hyperbolic discounting (a detailed discussion of the subject follows in ch. 3)

is implied if $f_{0,T/2} < f_{0,T} < f_{T/2,T}$ (2), in order that $f_{0,T} = f_{0,T/2} \times f_{T/2,T} = f_{0,T/2,T}$[7]. Now, (1) holds either if $f_{0,T/2} = f_{T/2,T} < f_{0,T}$ (3) or if (2) holds; hence subadditive discounting may (in the first case) or may not (in the second one) occur in addition to hyperbolic discounting, however (3) also leads to the observation of declining implied discount rates ($f_{0,T/2} < f_{0,T}$).

Table 1

Median responses and (continuously compounded discount rates in percent).

	Amount of early prize	Later prize paid in		
		3 mo.	1 yr.	3 yrs.
(A)	$15	$ 30 (277)	$ 60 (139)	$ 100 (63)
	$250	$ 300 (73)	$ 350 (34)	$ 500 (23)
	$3000	$3500 (62)	$4000 (29)	$ 6000 (23)
		6 mo.	1 yr.	5 yr.
(B)	$75	$ 100 (58)	$ 200 (98)	$ 500 (38)
	$250	$ 300 (36)	$ 500 (69)	$ 1000 (28)
	$1200	$1500 (45)	$2400 (69)	$ 5000 (29)
		1 mo.	1 yr.	10 yrs.
(C)	$15	$ 20 (345)	$ 50 (120)	$ 100 (19)
	$250	$ 300 (219)	$ 400 (120)	$ 1000 (19)
	$3000	$3100 (39)	$ 400 (29)	$ 10000 (12)
	Amount of early fine	Later fine due in		
		3 mo.	1 yr.	3 yrs.
(D)	$15	$ 16 (26)	$ 20 (29)	$ 28 (20)
	$100	$ 102 (6)	$ 118 (16)	$ 155 (15)
	$250	$ 251 (1)	$ 270 (8)	$ 310 (7)

The reason for that observation is that in the class of experiments such as the one described in Thaler (1981) the effects of delay (how far in the future the payment is to be received) and interval (the time between the two payments) are confounded. While hyperbolic discounting regards the delay effect, subadditive discounting accounts for the interval effect. In an attempt to disentangle those effects, Read designs three

[7] According to hyperbolic discounting, the discount factor over an interval is not altered by its partition on smaller subintervals. Nevertheless, a decline in discount rates is predicted.

experiments in order to test for the existence of subadditive (hypothesis H1: $f_{T\cdot 1} > f_{T\cdot 3}$) and hyperbolic (hypothesis H2: $f_{T\cdot 3\cdot 1} < f_{T\cdot 3\cdot 2} < f_{T\cdot 3\cdot 3}$) discounting[8], where $f_{T\cdot n\cdot i}$ is the discount fraction of the ith subinterval for a payment delay from now to T divided into n equal subintervals, and $f_{T\cdot n}$ is the respective compounded discount rate. In these experiments, T=24 months and n equals 1 or 3.

As it can be seen from Table 2, Read finds that the discount rate for a 24-month time interval is lower than the compounded discount rates of its three constituent 8-month subintervals, while the discount rates of the three 8-month intervals are roughly equal to one another, if not lower for earlier intervals. Hence these findings are consistent with subadditive, albeit not with hyperbolic, discounting. Especially the last result (a lack of declining impatience) is troubling, suggesting that in the previous experiments it was observed due to the confounding between the delay and interval effects. In any case, declining impatience can be explained equally well, if not better, by subadditive discounting. However, subadditive discounting cannot account for the 'time preference reversal' effect described earlier. Moreover, the lack of a considerable amount of studies verifying these results should make us reluctant to fully embrace subadditive as a replacement of hyperbolic discounting.

Table 2

Median discount fractions for all conditions of experiments 1-3[a]

Timing	Discount fraction				
	$f_{T\cdot 1}$	$f_{T\cdot 3}$	$f_{T\cdot 3\cdot 1}$	$f_{T\cdot 3\cdot 2}$	$f_{T\cdot 3\cdot 3}$
Exp 1					
SS	0.47	0.39	0.74	0.68	0.67
LL	0.60	0.39	0.81	0.71	0.72
Exp 2					
SS	0.43	0.32	0.68	0.71	0.64
LL	0.56	0.28	0.72	0.64	0.67
Exp 3					
SS	0.62	0.36	0.73	0.68	0.64
LL	0.61	0.36	0.72	0.68	0.68

[a] H1: $f_{T\cdot 1} > f_{T\cdot 3}$; H2: $f_{T\cdot 3\cdot 1} > f_{T\cdot 3\cdot 2} > f_{T\cdot 3\cdot 3}$; H4: $f_{T\cdot 1} > f_{T\cdot 3\cdot 1}$

[8] Also H4: $f_{T\cdot 1} < f_{T\cdot 3\cdot 1}$ (positive time preference) is employed as a consistency check.

On the other hand, Rubinstein (2003) conducts an experiment that contradicts the predictions made by hyperbolic discounting. Rubinstein presents the subjects with an experiment describing two similar situations (the conditions of the second situation are in parentheses): They are about to receive a new stereo system in replacement of their current one in 60 days (tomorrow), having to pay $960 ($1080) upon receipt. The question is whether they are willing to postpone delivery and payment by one day (60 days), for a discount of $2 ($120). The answers expected from the hyperbolic discounting hypothesis must be similar in both questions; individuals not willing to wait one day for $2, 60 days from now, should also not be willing to wait one day for $2 at any earlier date (actually they would require higher discounts, otherwise discount rates would not be declining, as hyperbolic discounting predicts). Hence, they should not be willing to postpone receipt 60 days from now for a $120 discount, by transitivity. However, Rubinstein reports that, from the 84 participants who answered both questions, about a quarter switched, choosing the one-day delay for $2, but not the 60-day delay for 120$[9].

The explanation Rubinstein provides to resolve this apparent contradiction is that maybe decision-makers simplify the choices they face by applying similarity relations to them, as Tversky (1977) proposed first. According to this approach, when facing a choice between dated prizes (x, t) and (y, s), decision makers use a procedure where they first check for similarities along the money and time dimensions, and look for dominance[10]. If dominance occurs in both dimensions, the decision is obvious. If not, the decision maker checks the two options for similarities along one dimension, and if this happens, he decides based on the other dimension. If this procedure does not give a definitive answer, he uses a different criterion. In the aforementioned experiment, if today we are at t=0, the first choice was between dated prizes $(-960{,}60)$ and $(-958{,}61)$, while the second one between dated prizes $(-1080{,}1)$ and $(-960{,}61)$[11]. Some subjects probably thought that in the first case there was a stronger similarity along the money compared to the time dimension (a $2 discount is not critical, while 60 days might differ more substantially from 61), so they chose not to delay receipt;

[9] 43% and 31% of the participants opted not to delay delivery in each of the two questions.
[10] Dominance in the money dimension is indicated by higher amounts (larger x's), and in the time dimension by dates closer to the future (lower t's).
[11] Negative values indicate that the amount is to be paid.

in the second case there was obviously no similarity in any dimension, so those subjects gave more weight to the money dimension and opted for delay.

As for the 'time preference reversal' observation, as described above, alternative explanations exist, the aforementioned procedure being among them. In the example provided by Ainslie & Haslam, it is possible that the subjects found delays of 6 and 8 years to be quite similar, thereby electing to wait for the larger check, but today and two years from now are not quite the same; so a different criterion is used for a decision to be reached. A drawback of this approach however is that the alternative criterion, as well as what exactly constitutes similarity, is not specified, thus it would be harder for this procedure to be modeled and used for predictions; instead, nearly every result could be rationalized by that vague notion of 'similarity' and an arbitrary decision criterion.

In addition, the 'immediacy effect'[12], describing the tendency to succumb to temptations for immediate pleasure, could account for 'time preference reversal'. The intuition behind this result is that we put a lot more weight on immediate gratification, compared to a delayed one; however, that happens only if the choice we are facing has instant ramifications. Considering the aforementioned example with the steak and the salad, the fact that the decision maker is following a diet that requires her to have salad for dinner, yet when going to a restaurant cannot avoid the temptation and orders a steak, need not imply that she is a hyperbolic discounter; it could rather be that the immediate nature of gratification made her lapse, and possibly this was not her original plan or a matter of sober intertemporal choice, as hyperbolic discounting implies[13]. Once more, the way the experiments are usually designed, like the one in Ainslie & Haslam, it is not possible for a distinction between the immediacy effect and hyperbolic discounting to be established[14]. The laboratory evidence asking subjects to compare sooner and later *future* rewards are scarce and inconclusive.

[12] See e.g. Hoch & Loewenstein (1991)

[13] In that case, whether preference reversal is observed depends on the time of delay (d), the amounts x_p and x_s, and the subsequent time t_s.

[14] Rewards are usually monetary, and, although not obvious, the immediacy effect might also occur for money; at least the available evidence does not contradict it, since in most experiments it is not implied that the amount will be paid *immediately*, and the subjects might have devalued it accordingly. In addition, money received right now could be used to purchase those goods that provide immediate satisfaction.

Despite the critiques presented above, which point out possible deficiencies underlying the hyperbolic discounting approach, nevertheless it has emerged as one of the best alternatives to the even more pronounced deficiencies of the DU model. This happens both because the proposed alternatives suffer from their own issues, such as the generality and vagueness of predictions of the ‘similarity relation procedure’, or because there are not sufficient evidence in support of an alternative like subadditive discounting.

2.3 Other deficiencies of the DU model

It is worth mentioning that declining discount rates are not only observed when delay intervals grow larger. The study mentioned earlier by Thaler (1981) demonstrates two other interesting findings that rebut the assumption of constant discounting, the ‘sign’ and the ‘magnitude’ effect. The ‘sign’ effect, otherwise termed ‘loss aversion’, as has been already presented, refers to the empirical observation that gains are discounted heavier than losses. To capture it, Thaler asked a fourth subset of participants to determine the amount they were willing to pay in order to postpone the payment of a traffic ticket (likewise to the other forms, three different fine amounts and delay intervals were considered). As can be seen from the comparison between the first and fourth forms of Table 1, where the respondents faced situations differing only on whether a prize or a fine was involved, the implied discount rates are much lower in the case of monetary losses compared to monetary gains. It seems that while an individual would require a compensation in the form of a considerably increased prize in order to postpone its receipt, she would be willing to pay only a marginally higher fine in order to delay payment, cultivating a ‘let’s get it over with’ attitude toward losses. In fact, based on the results presented for losses, one could hardly argue even for a declining discount rate, at least for this study.

As for the ‘magnitude effect’, which refers to the discounting of smaller monetary outcomes at a higher rate compared to larger ones, one can handily see, from form C of the Table 1, that the median implicit discount rate required by the subjects, in order to postpone by one month a $15 prize, was 345 percent, while for a similar delay period of one month and a $3000 prize, the required implicit discount rate has fallen sharply at about 40 percent. It is also worth mentioning that a similar pattern of discount rates declining with the size of the amount involved, if the delay window is held constant, happens in form D and in the case of losses, albeit to a much smaller extent, providing further substantiation to the proposition that losses differ fundamentally from gains. One possible explanation for this pattern, according to Thaler, is that there is a fixed psychic cost associated with waiting for a reward, which requires some mental effort[15]. This explanation can also account for the aforementioned pattern of discount rates decreasing with the length of time. To put it another way, if the mental effort required for waiting is not proportionally linked to the amount of the reward or the waiting time, then a part of the corresponding cost will be fixed and declining discount rates will be observed. It is needless to say that analogous results have been replicated by many other studies.

Although the relevant literature is mostly occupied with the assumption of constant discounting, that by no means suggests that the other assumptions of the DU model are immune to critique. First of all, many researchers have demonstrated that people generally prefer increasing sequences of outcomes, instead of decreasing ones, even if their total value is significantly lower. This is contrary to the predictions made by the DU model, since for positive time preference, the optimal distribution of outcomes that are spread over time would prescribe the best outcome to be placed at the beginning, with gradually worsening outcomes following, until the worst outcome is reached and placed last. Thus a decreasing sequence will be created; the reason being that positive discounting would make the most undesirable results feel less bad in the distant future, at least from today's perspective. For instance, Loewenstein & Sicherman (1991) found that an increasing wage profile was preferred by a majority of subjects over a declining or a flat one, yielding the same total value, for an identical job, even when they were reminded of the possibility that the last two options would produce a dominating consumption stream if invested properly.

Possible explanations for preference for improvement include the savoring and dread effect (Loewenstein, 1987), along with loss aversion and adaptation (Kahneman & Tversky, 1979). According to the first effect, individuals derive utility not only from consuming, but also from anticipating consumption of a good[16]. Hence improving

[15] That way the problem of intertemporal choice is described in terms of a self-control problem.
[16] The idea of utility from anticipatory consumption on Loewenstein's article (1987) incorporates negative utility from undesirable consumption.

sequences allow the best outcomes to be savored at the end, after a build-up of anticipation, which increases the utility derived from them, while dread is minimized by getting over with undesirable outcomes quickly, preventing accumulation of negative utility due to waiting. Loss aversion describes the greater emphasis put on losses, compared to gains, by individuals, compared to a reference point, which is usually determined by more recent outcomes, and adapted according to them. In that case, changes, rather than levels, of say consumption will be the determinants of value. Hence an improving sequence will be perceived as a series of gains from the relevant reference point, while a declining one will provide a series of relative losses, the former option being much more attractive in the context of loss aversion. Although the first effect applies to single outcomes as well as sequences, loss aversion is related only to outcome sequences. Consequently, it is expected that a more pronounced desire for improvement will be associated to sequences compared to single outcomes, since two effects operate instead of one.

Table 3

1. *Which would you prefer if both were free?*	n = 95
A. Dinner at a fancy French restaurant	86%
B. Dinner at a local Greek restaurant	14%
For those who prefer French:	n = 82
2. *Which would you prefer?*	
C. Dinner at the French restaurant on Friday in 1 month	80%
D. Dinner at the French restaurant on Friday in 2 months	20%
3. *Which would you prefer?*	
E. Dinner at the French restaurant on Friday in 1 month and dinner at the Greek restaurant on Friday in 2 months	n = 82 43%
F. Dinner at the Greek restaurant on Friday in 1 month and dinner at the French restaurant on Friday in 2 months	57%

Loewenstein & Prelec (1993) designed an experiment to confirm that hypothesis, asking undergraduate students three questions, with the results presented in table 3. When subjects were presented with their preferred alternative according to Question 1, the French dinner, as a single outcome, they overwhelmingly preferred immediate enjoyment, as the responses to Question 2 indicate. Nevertheless, when it was presented as part of a sequence also involving Greek dinner, a majority preferred to delay it, confirming our expectation for increased desire for deference of the most preferred alternative. It is worth mentioning, however, that a considerable number (20%) of the subjects opted to delay the French dinner even when it was presented as a single outcome.

Ample empirical evidence also exist against the consumption independence assumption of the DU model. Once again, Loewenstein & Prelec (1993) provide us with an interesting experiment, their population consisting of museum vistors. As can be seen in the following table, the only difference when choosing between the sets of consumption profiles prescribed by options A, B and C, D is their fifth weekend dinner plan. Thus, an agent preferring option A to B should also prefer C to D, since the substitution (across both choice sets) of the eat-at-home dinner with a fancy lobster dinner at the fifth weekend should make no difference. In any other case, additive separability[17], and consequently consumption independence, would be violated. Nevertheless, this was exactly what happened, with an overwhelming majority of subjects choosing option B over A, but with options C and D being virtually tied. The most plausible explanation is that subjects have a preference for spreading consumption over time. When the French dinner is the only fancy one in the profile, subjects almost unanimously prefer to place it in the middle of the interval. However, when the profile contains a second fancy dinner, they weakly prefer to spread these fancy dinners over the whole interval. It is interesting to note that neither habit formation nor loss aversion models can capture this preference for spread consumption, although they could capture a preference for improving sequences.

[17] Additive separability occurs when altering the elements of two sets in the same way does not lead to a preference reversal.

Imagine that over the next five weekends you must decide how to spend your Saturday nights. From each pair of sequences of dinners below, circle the one you would prefer. "Fancy French" refers to a dinner at a fancy French restaurant. "Fancy Lobster" refers to an exquisite lobster dinner at a four-star restaurant. Ignore scheduling considerations (e.g., your current plans).

first weekend	second weekend	third weekend	fourth weekend	fifth weekend	
		Option A			
Fancy French	Eat at home	Eat at home	Eat at home	Eat at home	[11%]
		Option B			
Eat at home	Eat at home	Fancy French	Eat at home	Eat at home	[89%]
		Option C			
Fancy French	Eat at home	Eat at home	Eat at home	Fancy Lobster	[49%]
		Option D			
Eat at home	Eat at home	Fancy French	Eat at home	Fancy Lobster	[51%]

All the evidence provided in this chapter point to the direction of abandoning the DU model as a benchmark for analyzing individuals' actions involving intertemporal choices. The reason is not only that the available empirical evidence about human behavior seem to refute the most critical assumptions of that model, but most importantly that these aberrations from DU predictions cannot be perceived as mistakes made by the individuals. As was mentioned earlier for example, even when subjects were told that a decreasing wage sequence permitted strictly higher consumption compared to an increasing one of equal nominal value, if subjects properly saved and invested their earnings, they upholded their initial decision. Whether this was because e.g. they were not confident that they would exercise the required self-control to maintain their initial consumption plans, or because they were wary of the implications of a declining sequence, the matter of fact is that individuals did not think that they should behave according to the predictions of the DU model; thus their choices should not be viewed as mistakes. This is a testament to the fact that the DU model bears neither descriptive nor normative utility, and a need for alternatives is imperative.

Chapter 3

Modeling time inconsistency

The fact that there exists ample experimental evidence contradicting the discounted utility model, as already discussed, has led economists to use different ways to model intertemporal choice. In most of the examples presented in the previous chapter, agents behaved in ways that departed from the predictions of the discounted utility model. The single most important deviation was the observation of declining discount rates, thereby leading to the questioning of the validity of an implication of constant discounting, namely time-consistency. As was argued in the previous chapter, constant discounting is a sufficient, albeit not necessary, condition for time-consistency. Hence, decision-makers that discount hyperbolically may have time-inconsistent preferences, presented in the 'time preference reversal' result described in the previous chapter.

3.1 The two extremes: Sophistication and naïvete

Strotz (1956) was the first to acknowledge this fact and formally model time-inconsistent agents. In his model, time-inconsistent agents, having to choose an optimal consumption plan, periodically deviate from the one initially chosen, which was viewed as optimal in an earlier date, since their time preferences have changed. To combat this tendency, an agent aware of his time-inconsistency, i.e. of the fact that his future selves have different tastes, might employ two different strategies to combat it: the 'strategy of precommitment', either precommiting herself to an, from the initial point of view, optimal plan of action irrevocably, or penalizing any future deviations from that plan; and the 'strategy of consistent planning', when precommitment is not feasible, where she chooses the optimal among the plans that she knows her future selves will actually follow.

Pollak (1968) elaborated Strotz's model, drawing an important distinction between agents that are aware of their inability to follow their initial plan (the type of agents Strotz was referring to), who he terms as 'sophisticated', and agents whose belief is that they will carry out their initial consumption path, only to revise it constantly after each period, who he terms as ‘naïve’. The relevant literature, for the most part, has adopted this distinction between sophisticated and naïve decision-makers when analyzing time-inconsistency. However, Strotz and Pollak neither used any specific functional forms to describe such preferences, nor they examined the implications of this distinction, only delineating the conditions characterizing the optimal plan.

3.2 Quasi-hyperbolic discounting: The (β,δ) model

The most widely adopted way by the literature to model hyperbolic discounting as the primary cause of time-inconsistent preferences is the (β,δ) model, popularized by Laibson (1997)[18]. In that paper, Laibson, in the context of a consumption-savings model as well, formulates the representative agent’s global utility from the perspective of period t as: $U_t = u(c_t) + \beta \sum_{i=1}^{T-t} \delta^i u(c_{t+i})$, where $u(c_t)$ is the instantaneous utility from consumption at time t, and β, δ are discount factors, hence $0 < \beta, \delta < 1$. The extra discount factor β captures the main implication of hyperbolic discounting, that immediate outcomes are discounted more heavily than distant ones: while the per-period discount rate between today and tomorrow is $\frac{1-\beta\delta}{\beta\delta}$, the discount rate between any two consecutive future periods is $\frac{1-\delta}{\delta} < \frac{1-\beta\delta}{\beta\delta}$ (since β < 1; obviously, when β = 1, this case is reduced to the standard DU model). However, Laibson terms this form of discounting ‘quasi-hyperbolic’, since the per-period discount rate, following a decline after the first period, remains constant afterwards, contrasting it to the constantly declining discount rates suggested by hyperbolic discounting[19].

Two are the main advantages of this functional form. First of all, it preserves the analytical tractability of the DU model by adding a constant (β) and slightly changing the discount factor. From the discussion of the previous chapter, it was noted that one

[18] Although Laibson’s work is recognized as the one upon which the relevant literature is based to model hyperbolic discounting, he was not the first one to use that specific functional form. It was actually Phelps & Pollak (1968), who used it to study intergenerational altruism in a consumption-savings model.

[19] A discount function that truly accounts for declining discount rates is $D(t) = (1+\alpha t)^{-\frac{\gamma}{\alpha}}$, $\alpha, \gamma > 0$, used by Loewenstein & Prelec (1992) among others. In that case, the per-period discount rate is defined as $-\frac{D'(t)}{D(t)} = \frac{\gamma}{1+\alpha t}$, which obviously declines with t. A higher α implies a function that departs heavily from constant discounting; however, in the limit as α goes to zero, this function reduces to the exponential discount function $D(t) = e^{-\gamma t}$, with γ being the per-period discount rate.

of the main reasons behind the widespread popularity of the DU model was precisely the facilitation of analytical derivations. However, quasi-hyperbolic discounting also accounts for the declining discount rates without hindering algebraic manipulations, as a truly hyperbolic discount function would do. The second, and perhaps most important, advantage of this discount function, compared to the general hyperbolic discount function, is that it also takes into account the already mentioned 'immediacy effect', i.e. the tendency to succumb to immediate pleasures, contrary to one's initial plans. That is because individuals who have preferences described by this function discount more heavily only outcomes actualized the following period, not ones that are more distant to the future; in essence they lapse only when the decision is right in front of them. That way their time inconsistency is handily revealed and they experience 'preference reversal', as they act in ways they will regret according to their previous preferences.

To illustrate that point, DellaVigna (2009) provides the following salient example: Consider a two period problem where an action leads to an immediate payoff b_1 in period 1 and a delayed payoff b_2 in period 2. That could either be an action regarding an investment good, in which case $b_1 < 0$ and $b_2 > 0$, or a leisure good, when $b_1 > 0$ and $b_2 < 0$ [20]. A decision-maker (DM) having (β,δ) preferences would *ex ante* (at t = 0) like to take the action if

$$\beta\delta b_1 + \beta\delta^2 b_2 \geq 0 \overset{\beta,\delta>0}{\Longleftrightarrow} b_1 + \delta b_2 \geq 0 \qquad (1).$$

However, come time t = 1, she will *ex post* decide to act if

$$b_1 + \beta\delta b_2 \geq 0 \qquad (2).$$ [21]

To incorporate the distinction between sophisticated and naïve agents, let $\hat{\beta}$ be the *belief* the DM has about her future preferences. Then $\hat{\beta} = 1$ corresponds to the naïve case, while $\hat{\beta} = \beta$ corresponds to the sophisticated case. Now the agent expects to consume in the future if

$$b_1 + \hat{\beta}\delta b_2 \geq 0 \qquad (3).$$

[20] Examples of the former include exercising or saving, where immediate cost or effort is followed by delayed pleasure, and of the latter a tempting food while on a diet, where an immediate pleasure leads to future cost.

[21] Obviously, if the DM is time-consistent (β = 1), she will indeed act as expected.

As an easy manipulation of (1) and (2) indicates, since $\beta < 1$, in the case of investment goods ($b_2 > 0$), if $b_1 \in [-\delta b_2, -\beta\delta b_2)$, the action is not taken despite the fact that it was desirable from an initial point of view, while in the case of leisure goods ($b_2 < 0$), if $b_1 \in [-\beta\delta b_2, -\delta b_2)$, the action is taken although it should not be chosen according to the initial preferences. Note that, in the case of sophisticates ($\hat{\beta} = \beta$), (3) reduces to (2), i.e. they act as expected, although it is still possible they will act contrary to their initial preferences, while in the case of naïves ($\hat{\beta} = 1$), (3) reduces to (1), i.e. they expect to act according to their preferences at time t, although for certain values of the parameters they will act in the opposite way.

Thus, by generalizing this result, quasi-hyperbolic discounting predicts inadequate consumption of investment goods and excessive consumption of leisure goods. In an effort to alleviate this self-control problem and restrict themselves from taking these non-optimal (from the t-self's point of view) actions, sophisticated agents will search for commitment devices that prevent them from consuming too little of investment goods and too much of leisure goods. That desire for commitment distinguishes them from naïve agents, who are unaware of their inability to abide by their initial plans, thereby deeming commitment as unnecessary.

An early example of how sophistication leads to the use of commitment mechanisms comes from Laibson's paper: The DM consists of a sequence of 'selves', that participate in a dynamic game, each one controlling the consumption decision for his respective period[22]. Each self, using illiquid assets as commitment mechanisms[23], try to constrain the consumption of future selves; or, to put it another way, the only reason that each self avoids overconsumption is that he faces a liquidity constraint endogenously imposed by his previous selves. Laibson using this model explains many empirical observations, such as the difference in MPC's for different types of assets, first noted by Thaler (1990), or the observed comovement between consumption and income. Because every self is endogenously liquidity constrained in the equilibrium path, his MPC equals one for liquid assets, because he wants to

[22] The way of modeling intertemporal choice as a game between a sequence of selves follows Strotz (1956) and Pollak (1968).
[23] Illiquid assets are usually modeled as imperfect commitment devices, either because their liquidation is lagging the receipt of the respective cash flow by one period, but can be used as a collateral for purchase of a liquid asset, as Laibson himself models them, or because there exist transaction costs, effectively penalizing their value if liquidated (see e.g. Angeletos et al.).

overconsume if allowed to, however it is close to zero for illiquid assets[24]. In addition, in a case of an unexpected increase in income at time t, the t-self is tempted to overconsume it, since he is constrained only for the expected part of his income; thus, unexpected increases in income will be followed by consumption increases.

Angeletos et al. (2001) provide further support for (β,δ) preferences in the context of consumption-savings models, where by calibrating models with (β,δ) and exponential preferences and comparing them, they find that the former can better approximate observations based on real-world data, such as that households hold only a slight share of their wealth in liquid assets, and have high credit card debt. The reason is that although both types of households equally dislike illiquid assets because they cannot be used for consumption smoothing in case of external shocks, households that discount hyperbolically at least value them as commitment devices. Hence, illiquid assets' net value is higher for that type of households, and they tend to hold relatively low liquid wealth. In addition, credit card debt is used at a higher rate than liquid assets to smooth future consumption, although it also comes at the cost of facilitating future lapses, because future selves might partly refrain from lapses due to the interest cost of borrowing. Both papers describe the use of illiquid assets as commitment devices employed by sophisticated agents to limit inadequate consumption of investment goods. However, no reason supporting the assumption of sophistication is provided, and an extension of the analysis in the case of naïvete would be interesting[25].

O'Donoghue & Rabin (1999), in a model studying procrastination (a deliberate delay in the execution of an onerous act), were the first to explicitly analyze the consequences of naïvete in the context of a model where agents have (β,δ) preferences. In their model, a DM has a finite amount of time ($T < +\infty$) to perform a task, and this task is associated with a reward and a cost schedule ($\boldsymbol{v} = (v_1,...,v_T)$ and $\boldsymbol{c} = (c_1,...,c_T)$, respectively). The DM's actions are described by a strategy $\boldsymbol{s} = (s_1,...,s_T)$, where $s_t \in \{Y, N\} \quad \forall t \in \{1,2,\dots,T\}$; s_t specifies whether she performs

[24] In the case of exponential discounting, income from liquid assets need not be consumed, but saved for the future, since all selves value future consumption the same way.

[25] As we will note later, behavioral evidence does not provide such a strong support in favor of sophistication. Part of the explanation for this choice might be that assuming naïvete is a stronger departure from standard economic assumptions; not only agents are time-inconsistent, but they also have 'non-rational expectations' about their preferences.

the task ($s_t = Y$) or waits ($s_t = N$) at time t under $\boldsymbol{s}$[26]. Also, since the task must be performed, it is required without loss of generality that $s_T = Y$ [27]. They study two separate cases: when the task is completed at time t, either the cost is incurred at time t and the reward comes at t+1 (the case of immediate costs, corresponding to investment goods from our previous analysis), or the reward comes at time t and the cost is incurred at t+1 (the case of immediate rewards, corresponding to leisure goods from our previous analysis).

The solution concept of perception-perfect strategy is constructed for the three types of agents: time-consistent (TC), naïfs [sic] (n) and sophisticates (s). That concept prescribes that in equilibrium, in every period the DM chooses the optimal action given her current preferences and her perception about her future behavior. Hence, in the case of TC's and naïfs, $\boldsymbol{s^{TC}}$ and $\boldsymbol{s^{n}}$ are defined as perception-perfect strategies if:

$$\forall t < T,\ \ s_t^k = Y \ \ iff \ \ U_t(t) \geq U_t(\tau) \ \ \forall \tau > t, \qquad k = TC, n$$.[28]

Essentially both TC's and naïfs follow the same decision process ('do the task now if it is optimal given your current preferences'), because naïfs falsely believe that they will behave as TC's if they wait, although they are present-biased[29]. Meanwhile sophisticates differ from naïfs in the perception of their future behavior, realizing their self-control problems; hence, their perception-perfect strategy $\mathbf{s^s}$ is defined as:

$$\forall t < T,\ s_t^s = Y \ \ iff \ \ U_t(t) \geq U_t(\tau') \ \ \text{where} \ \ \tau' \coloneqq \min_{\tau > t} \ \{\tau | s_\tau^s = Y\}$$

The above condition states that sophisticates will execute the task if it is optimal from their current perspective, after taking into account when their future selves will do it if they now choose to wait. Also, the period that agent-type *k* actually performs the task given their strategy $\boldsymbol{s^k}$ is defined as:

$$\tau_k \coloneqq \min_t \ \{t \big| s_t^k = Y\}, \qquad k = TC, n, s$$

O'Donoghue & Rabin show that in the case of immediate costs, $\tau_n \geq \tau_{TC}$, while in the case of immediate rewards, $\tau_n \leq \tau_{TC}$. This is a manifestation of the *present-bias effect*, which leads present-biased individuals to 'procrastinate' (wait when they

[26] As usual, $\boldsymbol{s}$ specifies what the DM does even in subhistories that will never occur, e.g. if $s_t = Y$, $s_{t'}$ is still specified for t' > t.
[27] A weaker, yet still equivalent, requirement would be: $s_T = Y$ if $s_t = N \ \forall t < T$.
[28] $U_t(j)$ denotes the global utility from self-t's perspective if the task is completed at time j.
[29] This is an alternative term that O'Donoghue & Rabin use to refer to hyperbolic discounters (β < 1)

should do it) when costs are immediate, while they 'preproperate' (do it when they should wait) when rewards are immediate; naïfs are solely influenced by this effect. However, sophisticates are also influenced by the *sophistication effect*, an acknowledgement from their part of their self-control problems, which leads them to alter their behavior now. The implication of this effect is that sophisticates complete the activity earlier than naïfs, regardless of the timing of costs and rewards; that is, $\tau_s \leq \tau_n$[30]. However, the interpretation given to this result is two-fold: in the immediate cost case, the sophistication effect alleviates their proclivity for procrastination; in the immediate reward case, it exacerbates their tendency for preproperation. Of course, whether sophisticates will end up better off than naïfs depends on the combination of both effects.

To illustrate their point, the authors give two examples (in both it is assumed that T = 4, β = ½ and δ = 1[31]): in the first, costs are immediate, rewards are received at time t = 5, and the reward and cost schedules are $\boldsymbol{v} = (\bar{v}, \bar{v}, \bar{v}, \bar{v})$ and $\boldsymbol{c}$ = *(3,5,8,13)* respectively. Then an easy calculation reveals that $\boldsymbol{s}^{\boldsymbol{TC}}$ = *(Y,Y,Y,Y)*, hence $\tau_{TC} = 1$, $\boldsymbol{s}^{\boldsymbol{n}}$ = *(N,N,N,Y)*, hence $\tau_n = 4$, and $\boldsymbol{s}^{\boldsymbol{s}}$ = *(N,Y,N,Y)*, hence $\tau_s = 2$. In the second example, rewards are immediate, costs are incurred at time t = 5, and the reward and cost schedules are $\boldsymbol{v}$ = *(3,5,8,13)* and $\boldsymbol{c} = (\bar{c}, \bar{c}, \bar{c}, \bar{c})$ respectively. Again an easy calculation reveals that $\boldsymbol{s}^{\boldsymbol{TC}}$ = *(N,N,N,Y)*, hence $\tau_{TC} = 4$, $\boldsymbol{s}^{\boldsymbol{n}}$ = *(N,N,Y,Y)*, hence $\tau_n = 3$, and $\boldsymbol{s}^{\boldsymbol{s}}$ = *(Y,Y,Y,Y)*, hence $\tau_s = 1$. The reasoning for the three groups are the following:

- TC's simply maximize the difference v_t - c_t in both cases, making always the optimal decision.
- Naïfs, in the first case, wrongly predict in each period that they will not procrastinate in the following one, so they underestimate the cost of procrastination now, leading them to perform the task in the worst possible time. In the second case, they believe that they have the willpower to wait until the last period, but finally give in to their self-control problem in the third period.
- Sophisticates end up better off than naïfs in the first case: they know that they will not do the task in the 3rd period, but they correctly predict that their

[30] Remember that sophisticates do it at time t if $U_t(t) \geq U_t(\tau')$, with τ' as defined previously, which is a weaker condition than the one required for naïfs: $U_t(t) \geq U_t(\tau)\ \forall t > \tau$.

[31] The ignorance of 'long-term discounting' is an innocuous assumption for comparison purposes.

period-2 self, being aware of period-3 self's self-control problems, will perform it; that way they afford to procrastinate in the 1st period. They do procrastinate, yet their sophistication mitigates that problem. In the second case however, they make the worst possible decision: knowing that her period-2 self will give in (since that self already knows that the period-3 self will give in), self 1 figures out that she might well do it now. So they preproperate worse than a naïf, due to their cognizance of their future selves' self-control problems. This result is a testament to the fact that sophistication does not always lead to increased welfare.

It is noteworthy that despite the fact that in both of the aforementioned examples the present-bias effect overtakes the sophistication effect, this need not be the case. On the contrary, it might well be that sophisticates preproperate even if the costs are immediate, in which case the sophistication effect is greater than the present-bias effect[32]. Another example from O'Donoghue & Rabin is telling: In a situation like the first one, with the only differences being that T = 3, the rewards are received a period later, $\boldsymbol{v}$ = *(12,18,18)* and $\boldsymbol{c}$ = *(3,8,13)*, it holds that $\tau_s = 1$, $\tau_{TC} = 2$ and $\tau_n = 3$. Now sophisticates preproperate, knowing that if they do not they will end up procrastinating, which is more costly. Although they still are better off than naïfs, the possibility that sophisticates may exert 'preemptive overcontrol' reminds us that both effects must be taken into account when analyzing sophisticates' behavior, in order to avoid situations like that, where the observed behavior seemingly contradicts present-biased preferences.

3.3 Partial naïvete: The $(\beta, \hat{\beta}, \delta)$ model

While a great part of the literature studied the implications of naïvete by assuming either completely naïve or completely sophisticate agents, these cases are somewhat extreme. There is no particular reason to assume that a person either is completely unaware about his self-control problems or that he correctly identifies their extent. It might well be that, although he understands their existence, he underestimates the extent of these self-control problems. O'Donoghue & Rabin (2001) were the first to

---

[32] As we already noted, the opposite is impossible: a sophisticate would never procrastinate when immediate rewards are involved, since then their present-bias is augmented by their sophistication.

diverge from that extreme dichotomy, making an interesting elaboration on the previous model by introducing the notion of *'partial naïvete'*. They define a partially naïve DM in a straightforward way: they let $\hat{\beta} \in (\beta, 1)$[33]; thus the (β,δ) model is modified as the *$(\beta, \hat{\beta}, \delta)$* model. This way we can examine what happens in between the two extreme cases.

In context of the example provided by DellaVigna (2009), for instance, from (2) and (3), a similar analysis to the previous predicts that: i) the case of investment goods ($b_2 > 0$), if $b_1 \in [-\hat{\beta}\delta b_2, -\beta\delta b_2)$, the action is not taken although it was expected to, and ii) in the case of leisure goods ($b_2 < 0$), if $b_1 \in [-\beta\delta b_2, -\hat{\beta}\delta b_2)$, the action is taken although it was not expected to. To generalize, in both cases, a partially naïve agent is overconfident about his ability to abide by his initial plans, overestimating his consumption of investment goods and underestimating his consumption of leisure goods (in contrast, a naïve agent is always overconfident, while a sophisticate one never is).

Partial naïvete seems a much more plausible assumption compared to the two extremes. One of the earliest studies providing evidence in favor of partial naïvete comes from Ariely & Wertenbroch (2002). They studied professionals enrolled in an executive education class at MIT, which required them to complete three coursework assignments, and asked them to specify a binding deadline for each (if they missed it, they would get lower grades). While a time-consistent student would choose the last day of the semester as the deadline[34], 68 percent of the deadlines were set before the last week, indicating a demand for commitment. The participants, possibly realizing their self-control problem, which would lead them to underconsume an investment good like homework completion, viewed early deadlines as a commitment device, and opted to use them to restrict their future-selves.

However, in a follow-up experiment where the participants (students having three homework assignments to complete within a 21-day period) were split in three groups: students in the control group could return the assignment at any time during

[33] Remember that naïfs are defined as having $\hat{\beta} = 1$, while sophisticates as having $\hat{\beta} = \beta$.

[34] Since they received no feedback even if they delivered at an earlier date, while at the same time they would have lower flexibility, it would not be optimal to set early deadlines, i.e. to restrict your choice set (when to do the assignments).

the specified period, students in the first treatment group could set their own deadlines, facing a problem much like the previous one, and finally students belonging in the second treatment group faced three equally-spaced deadlines of one week each. While the first treatment group performed significantly better than the control group, corroborating the importance of deadlines as commitment mechanisms, the second group did significantly better compared to both other groups, suggesting that self-imposed deadlines were not set optimally. The main implication of these findings is that although the participants recognized their self-control problems searching for commitment devices, they underestimated these problems, which led them not to choose these commitment devices optimally, as would be expected from a partially naïve agent.

O'Donoghue & Rabin (2001) examined the implications of assuming partially naïve agents in the context of their previous work on procrastination. They slightly change the framework on which they perform their analysis: now the time is infinite, and only the case of immediate costs is analyzed (thus the results mainly apply to investment goods). Moreover, the DM firstly chooses from a menu of tasks which one to complete (if she decides to complete any at all), and then decides when to complete her chosen task (however she is not required to). Hence, the possibility of infinite procrastination arises. Although a short delay in completion is possible, for the most part it is not particularly inefficient; for that reason, the authors regard that the DM procrastinates only if she does not complete any task at all, provided that there is a task worth completing.

In the case of one task for example, the DM might procrastinate infinitely if her maximum tolerable delay $d*$ (indicating that she prefers to complete the task today than complete it after $d \geq d^* + 1$ periods) is greater than zero, depending on her perception about her future behavior. For a sophisticate DM, the authors show that any perception-perfect strategy will lead her eventually to complete a task worth doing, so she will never procrastinate according to the new definition[35]. For a partially naïve DM however, whose perceived future maximum tolerable delay is $\widehat{d^*}$, the possibility of infinite procrastination arises if $\widehat{d^*} + 1 \leq d^*$. Because the DM, if she has to choose between doing it now or after $\widehat{d^*}$ periods, will always believe that she

[35] Such a strategy is the one prescribing to do the task every $d^* + 1$ periods, so that the DM will reach a point that she prefers to complete the task then.

can wait for now, yet the task will eventually be done, she will postpone its completion indefinitely. The authors extend this logic to the case of a menu of tasks in an analogous manner, showing that even a small degree of naïvete suffices for procrastination and welfare loss compared to the full sophistication case in some environments, as was previously argued by O'Donoghue & Rabin (1999), who came to the same conclusions for any case when comparing full naïfs and full sophisticates with immediate costs involved. Thus, we could say that partially naïve agents behave in between fully sophisticated and fully naïve ones.

Beyond extending the previous model to incorporate partial naïvete, O'Donoghue & Rabin also find two interesting new results. The first finding is that providing additional options may in fact induce procrastination, although if they were not available the DM would not procrastinate, violating the axiom of independence of irrelevant alternatives. That can happen when the new task has higher long-term benefits than those already available. In that case, the DM will ignore them and decide to perform the new task; but if its cost is relatively high compared to its benefit, she will procrastinate indefinitely, despite the fact that she might have not procrastinated if she has chosen one of the initial options. This result holds again for any degree of naïvete, and it can be shown that such a procrastination-inducing task always exists.

The second finding is that 'increasing importance can exacerbate procrastination'. Again the reason is that more important goals are worth pursuing, even when they come at a higher cost; yet this increased cost might lead even a slightly naïve person to procrastinate. Of course increasing importance in the one 'task' case mitigates procrastination, so it mitigates it for any specific option in the 'menu of tasks' case; but when multiple options are available, costlier tasks become more attractive (because costs are underestimated), and the fact they are costlier increases the likelihood of procrastination, as was explained previously. These observations can explain many observed patterns, including retirement planning, such as why people often delay transferring savings from checking to higher-interest accounts; since that decision is of paramount importance, procrastination probably plays a part.

Another application of partial naïvete with $(\beta, \hat{\beta}, \delta)$ preferences, in a market setting this time, is done by DellaVigna & Malmendier (2004). They develop a model of optimal pricing with two-part tariffs, consisting of a lump-sum fee and a per-usage price where time-inconsistent agents face a monopolistic firm[36]. In the case of investment goods, the monopolist charges time-consistent agents with his marginal cost per usage, but charges with a per-usage price below marginal cost time-inconsistent ones[37]. The reason for the difference in the optimal pricing scheme is two-fold: a commitment and an overconfidence effect taking place simultaneously. A consumer at least partially *aware* of his self-control problem perceives the lower per-usage price as a commitment device to mitigate his self-control problem (in that case inadequate consumption). At the same time, a consumer at least partially *overconfident* about his self-control problem can be exploited by the firm for his misperception of his future behavior. In essence, the firm lures the consumer with the low per-usage price, which she thinks she can take advantage of, compensated with an increase in the lump-sum fee.

An analogous result can be obtained in the case of leisure goods: now the optimal per-usage price is above the marginal cost for time-inconsistent agents, if they are partially aware of that fact they want to limit overconsumption, and because they believe that underestimate their consumption if they are partially naïve, they are lured by a low lump-sum fee followed by a high per-usage price. These results hold even if competitive markets are assumed, they only difference being in the equilibrium fee: instead of setting it to extract the consumers' surplus, firms set it equal to their fixed cost, if any, in order to make zero profits in equilibrium; yet they depart from marginal-cost pricing if they face time-inconsistent agents in the ways delineated above.

Regarding the welfare effects of time-inconsistency, consumer welfare (both in monopolistic and competitive markets), firm profits (in the case of monopoly) and total welfare are not affected by the degree of time-inconsistency $(1 - \beta)$ if consumers are fully sophisticated, since then the overconfidence effect is negated. Nevertheless, partially naïve consumers have lower surplus compared to sophisticates in both market environments, and the monopolist's profits are strictly increasing to their degree of naïvete $(\hat{\beta} - \beta)$; in the end total welfare is reduced under naïvete. The reason is that a monopolist deviates from the first-best outcome, induced by marginal-

[36] Most models in the Behavioral Industrial Organization field assume time-inconsistency (or bounded rationality in general) only on the demand side; firms are rarely assumed to be boundedly rational.
[37] Obviously, the lump-sum fee is set so as to make the participation constraint of the consumers binding for both types of agents.

cost pricing, in order to appropriate a larger part of the consumer's surplus, while in a perfectly competitive market the first-best is still not achieved, however the firm now is not the beneficiary of the surplus, thus the consumers' welfare is not reduced so heavily. Consequently, the loss of consumer welfare (the difference in consumer welfare between monopoly and perfect competition) is increasing to their degree of naïvete.

Examples from the health club industry (DellaVigna & Malmendier (2006)) and the credit card market (Ausubel (1991)) provide strong evidence in favor of partial naïvete on both investment and leisure goods. DellaVigna & Malmendier (2006) use data from three U.S. health clubs that offer two types of contracts: one monthly contract with lump-sum fee of $70 on average or a pay-per-visit scheme of about $10 per visit. One would expect them to choose the monthly contract only if they expect to attend at least 7 times per month, so that their price per expected visit is lower than the price per visit. The authors find that customers who chose the monthly contract that, on average, expected to visit the gym 9.5 times per month; nevertheless, they end up attending the gym only 4.17 times per month, with an average price per attendance of over $17 (only 20 percent pay an average price of less than $10, making their choice worthwhile). The authors consider many alternative explanations, but they conclude that the most plausible one is that attendees are partially naïve. Since exercising is an investment good with immediate costs and deferred benefits, we expect partially naïve agents to sign up for monthly contracts both as a commitment device (that way they lower the cost per attendance from $10 to $0, hoping that this will suffice to constrain their future selves) and because they overestimate their consumption (health club attendance), and in the end to underconsume the good. However, the last reason seems to play a bigger part, and health clubs take advantage of consumers' naïvete making higher profits (the average attendee foregoes about $600 in his membership due to the higher cost per attendance).

Ausubel (1991) observes that credit card interest rates remained constant during the 80's, despite the presence of almost 4000 U.S. banks, and a variation in their marginal costs. This finding suggests the bank credit card market seems to depart from the predictions of competition, with above-marginal-cost pricing. Ausubel attributes it to the overconfidence of consumers about credit card borrowing, which is a leisure good (it allows you to increase current consumption at the expense of future one). Because consumers believe they can exert self-control, borrowing wisely and repaying their debt within the specified 30-day period without interest, banks induce them to sign-up by not requiring an initial fee (some even provide benefits), but require high interest rates, pricing above marginal cost (the interest rate is the equivalent of per-usage price of credit card debt). In addition, the fact that almost three quarters of consumers pay finance charges on their outstanding balances seems to support the consumer naïvete assumption. These findings are in line with the predictions of DellaVigna & Malmendier (2004) for leisure goods.

3.4 Alternatives to partial naïvete: Frequency and magnitude naïvete

Although $(\beta, \hat{\beta}, \delta)$ preferences are by far the most common way to model partial naïvete, let us turn in this part to two alternative ways that have been proposed by the literature. These are magnitude naïvete (Spiegler (2011)), or 'projection bias' as it was termed by Loewenstein, O'Donoghue & Rabin (2003) who introduced the notion, and frequency naïvete introduced by Eliaz & Spiegler (2006). An individual's preferences exhibit projection bias if their projected instantaneous utility from consumption $u(\widetilde{c,s'})$ lies between her true future utility $u(c,s)$ and her current utility $u(c,s')$, where s' is her current state[38] and s her future one; or more formally, if $u(\widetilde{c,s'}) = \alpha u(c,s') + (1-a)u(c,s)$, where $a \in [0,1]$ measures the 'magnitude' of naïvete, hence the alternative name. $\alpha = 0$ corresponds to the full sophistication case (the DM correctly predicts her future utility), and $\alpha = 1$ to the full naïvete case (the DM fails to perceive any change in her utility), with the intermediate values implying partial naïvete, and a higher α implies a stronger projection bias. Once again, even the slightest degree of naïvete can lead to time-inconsistent preferences (the DM fails to exactly pin her future instantaneous utility, and consequently revises her initial plans after realizing that they are not optimal); however, a fully sophisticated agent without projection bias differs from a fully sophisticated with (β,δ) preferences in the sense that the former is not time-inconsistent, since in that case time-inconsistency arises

[38] An individual's state could depend on her previous consumption, as in a habit formation model (see e.g. Pollak (1970)), or any other factor that can influence the experience of consumption, such as her mood (see e.g. the notion of visceral influences, due to Loewenstein (1996)).

from misprediction of future utility, not by a preference for immediate compared to delayed gratification.

In the case of frequency naïvete, which is analyzed in a two-period model, the DM is assumed to assign probability $\theta \in [0,1]$ to the event that her period-2 self's utility will be the same as the period-1 self's *(u)*, although in fact it will be given by the function *v*. The analysis is the same as with frequency naïvete, with θ playing the role of α, however the source of time-inconsistency in that case is not a misprediction of how much preferences change; it is rather an underestimation of the possibility of that change[39]. Hence, it constitutes a more substantial departure from $(\beta, \hat{\beta}, \delta)$ preferences.

It is worth noting that which type of partial naïvete one assumes may yield completely different predictions. An interesting comparison of the two cases comes from Spiegler (2011). In a setup where a monopolist faces a continuum of agents with differing degrees of naïvete, he offers two price schemes: one seen as a commitment device, aimed at sophisticates, and one exploitative, aimed at naifs. That way the monopolist is essentially screening agents according to their beliefs. However, the implications are radically different in the two cases: in the magnitude naïvete case, in equilibrium one price scheme is aimed at fully sophisticate consumers, while the other is used to pool all the other consumers, so that even the slightest degree of naïvete makes those consumers indistinguishable. In the frequency naïvete case, there exists a cutoff type of consumer, where more sophisticated and more naïve consumers than him are pooled together in the respective price schemes. The main reason for that difference is the respective market setting. In the first case, the monopolist knows the type of the consumer, thus he is able to exploit him even he is slightly naïve (this is a close analogue to the DellaVigna & Malmendier model of two-part tariffs, where the monopolist departs from marginal-cost pricing for all naïve agents, irrespective of their degree of naïvete). In the second case, the monopolist cannot know each consumer's type, so he cannot extract fully their surplus when they are sufficiently sophisticated, pooling them with full sophisticates.

These models also differ in other aspects. For instance, under magnitude naïvete increased sophistication might lead to decreased welfare. The monopolist may offer a

[39] An example would be a person on a diet underestimating the temptations that will lead him to lose his willpower and break his diet, not an underestimation of his willpower per se (as the other two models imply).

sufficiently expensive contract that is perceived as a commitment device, leading a more sophisticated agent to select it and end up worse because she still does not commit herself, having underestimated her self-control problem, whereas a fully naïve person does not even search for a commitment device (this possibility in fact is reminiscent of the sophistication-exacerbates-procrastination finding). In the context of frequency naïvete, increased sophistication at least does not decrease welfare.

Throughout this chapter, many examples have been presented of the way time-inconsistent agents can be modeled, and most of them can explain 'stylized facts' or observed patterns in many market settings, in ways that traditional rationality-based models cannot. However, a note of caution with regard to the last two examples: an important assumption in modeling time-inconsistency and naïvete is not only its degree (full, partial or none at all), but also its source (whether the agents simply underestimate their preference change, or it is not known whether this change will happen), since different implications follow from those assumptions.

Chapter 4

Implications of time-inconsistency: The case of addiction

The main intent of this chapter is to analyze the disparities that would emerge, depending on whether an agent was time consistent or not, in hers consumption choice of an addictive good. Through this comparison, the appropriateness of each model to describe addictive behavior is evaluated, as well as the difference in policy prescriptions that each perspective entails. This is done by analyzing the empirical evidence available for smoking consumption and sales; however, the results obtained could probably be used to describe other addictive behaviors without much refinement.

4.1 Rational addiction

Addictive behavior is closely related to the idea of habit formation (first introduced by Duesenberry (1949), elaborated by Pollak (1970), among others). In these models, an agent’s current utility depended on the present, as well as the past, consumption of a good, an effect that could readily describe addiction, or at least its reinforcement aspect (the fact that the higher the past consumption of an addictive good, the greater the desire for current consumption). However, many experimental studies[40] have found a second characteristic of addictive behavior, tolerance (the fact that the same level of consumption yields lower utility for higher levels of past consumption), that habit formation fails to explain.

Becker & Murphy (1988), in a groundbreaking article, employed a rational choice framework to analyze addiction, arguing that while many aspects of addiction seem irrational at first, they could be properly understood as rational choices of agents with

[40] E.g. Donegan et al. (1983)

stable preferences and optimizing behavior, who recognize the addictive nature of the good they consume.

They assume that an individual's instantaneous utility is of the form u(t)= u(y(t), c(t), S(t)), where y, c and S are the current consumption of the non-addictive good, the current and the past consumption of the addictive good, respectively. A utility function that is additively separable and strictly concave at its arguments is assumed. Past consumption resembles the stock of 'consumption capital', and is accumulated through a standard investment function, with a depreciation rate of δ.

Hence the problem of an agent with lifespan T and rate of time preference σ is:

max $\int_0^T e^{-\sigma t} u(y(t), c(t), S(t))dt$

s.t. $\int_0^T e^{-rt} \, [y(t) + p(t)c(t)]dt \leq A_0 + \int_0^T e^{-rt} w(S(t))dt$

$\dot{S} = c(t) - \delta S(t)$,

where the non-addictive good is treated as the numeraire, the initial value of assets and of consumption capital A_0 and S_0 are given, and w(·) is the earnings function, which is concave on the stock of consumption capital. The first-order conditions of this problem are:

$u_y(t) = \mu e^{(\sigma - r)t}$

$u_c(t) = \mu p(t) e^{(\sigma - r)t} - a(t) = \Pi_c(t)$ (1),

where $= \frac{\partial V}{\partial Ao}$, with V being the value function, and

$a(t) = \int_t^T e^{-(\sigma+\delta)(\tau - t)} \, u_s d\tau + \mu \int_t^T e^{-(r+\delta)(\tau - t)} \, w_s d\tau$ (2).

This expression measures the discounted utility and monetary cost (or benefit) of an additional unit of c, through its effect on the future stock. It essentially is the shadow price of an additional unit of stock. So the addictive goods can be categorized as harmful (if u_s, w_s < 0) or beneficial (if u_s, w_s > 0), depending on its effects on future utility and earnings. The second first-order condition points out that a rational addict chooses her consumption path of the addictive good, so that its marginal utility, u_c(t), equals its full price, Π_c(t), i.e. the sum of its (discounted) market price and the (discounted) money value of its effect on future consumption. Hence, for example, a rational addict consumes more of a beneficial compared to a harmful addictive good, ceteris paribus, recognizing its benefit.[41]

Becker & Murphy continue by assuming a quadratic functional form for u and w, and a constant price ($p(t) = p \;\; \forall t$), in order to analyze its dynamic implications. This leads to a quadratic value function[42] of the form

$$F(t) = a_c c(t) + a_S S(t) + \frac{a_{cc}}{2}[c(t)]^2 + \frac{a_{SS}}{2}[S(t)]^2 + a_{cs}c(t)S(t) - \mu pc(t),$$

with $a_{cc}, a_{ss} < 0$, due to the concavity of u and w.

They then proceed by giving the definition of addiction in this framework. The notion of adjacent complementarity is central to their analysis of addictive behavior.[43] An agent is addicted to a good iff his consumption pattern exhibits adjacent complementarity, i.e. iff a rise in his current consumption raises his future consumption as well (or to put it differently, iff his consumption decisions at different points in time are complements). They point out that while that happens when the marginal utility of c increases over time when the stock of consumption capital increases over time ($a_{cs} > 0$), this is not a sufficient condition for adjacent complementarity. The reason is that concurrently the full price of consumption, $\Pi_c(t)$, also increases over time[44], and that might lead to a fall in consumption if its increase rate exceeded the increase rate of marginal utility. A stronger condition that ensures the presence of adjacent complementarity is the following: $(\sigma + 2\delta)a_{cs} > -a_{ss}$ (3).

This inequality implies that if F is more concave in S (a_{ss} is higher in absolute value), the individual discounts the future less heavily (σ is lower), or the depreciation rate of consumption capital (δ) is lower, the full price of consumption increases, and if these factors are sufficiently strong, an increase in current consumption would lead to a fall in the future one.

It should be noted that this framework allows us to explain multiple aspects of addictive behavior. First of all, reinforcement is captured by the presence of adjacent complementarity, while tolerance can also be accounted for, since the same level of

---

[41] This happens since in the latter case a(t) is negative, and the full price and marginal utility of the addictive good are higher, leading to a lower c(t), since u(t) is concave at c(t).
[42] After substituting out for y(t) from its FOC.
[43] The following analysis holds both for harmful and beneficial addictions.
[44] Since $a_{SS} < 0$, so the accumulation of consumption capital leads to higher cost of future consumption in terms of decreased maximal utility that could be achieved.

current consumption yields lower utility for greater levels of past consumption, at least for harmful addictions.[45] However, we should note that this framework cannot explain the tolerance effect in the case of rational beneficial addictions.

Moreover, this framework can explain the bimodal distribution of consumption, usually associated with addictive goods. This happens through the possibility of an unstable steady-state, if adjacent complementarity is sufficiently strong. For example, if a cubic term of S with negative sign was added to the value function F, then this would lead to two steady-states, one stable and one unstable, with the consumption being higher in the stable steady-state.[46] Hence, if S_0 is higher than a particular threshold, consumption converges to the sizable stable steady-state level, while if it is lower compared to that threshold, it diverges to zero, and the individual selects to abstain. This explains the fact that consumption of highly addictive goods, such as cigarettes, is usually concentrated around two modes, one with sizeable consumption and one with infrequent use, near abstention. Unstable steady-states can also help explain 'binges', rapid increases of consumption within a short time period.

Another interesting result can be obtained from examining the role of the rate of time preference σ in the consumption decisions of the rational addict. While for most goods an increase in σ, i.e. heavier discount of the future, relative to the interest rate r is associated with substitution of future for current consumption, we can see from (1) and (2) that if the good is sufficiently addictive its full price might fall and its consumption might increase over time. Hence it is possible to obtain the otherwise counterintuitive result that individuals who care less about the future might increase the consumption of the addictive good constantly, due to the reinforcement aspect of addiction outweighing the heavy future discount.

Note that this model of rational addiction can also account for myopic behavior in the consumption of addictive goods. Myopic behavior occurs as σ gets larger and approaches infinity. Then the present value of the cost of an increase in the stock of consumption, $a(t)$, goes to zero,[47] and an agent can ignore the future effects of his current actions. As equation 3 implies, if $a_{cs} > 0$, then even a fully myopic agent can

---

[45] Remember that for harmful addictions, it holds that $u_s < 0$, which is the definition of the tolerance effect.

[46] Due to the fact that adjacent complementarity decreases with S, since a_{SS} gets more negative with higher S (from the negative sign of the term S^3)

[47] Assuming that σ=r.

be addicted ‘rationally’. However, whether an agent that discards completely from consideration any future effect of his behavior can be described as rational is a valid question. For example, due to adjacent complementarity, an anticipated future increase in the price of the addictive good will lead to a reduction in current consumption by a rational addict as a response to lower expected future consumption, but a myopic addict would not react to such a change.

4.2 Empirical findings on cigarette consumption

Many empirical tests have been designed to test for myopic vs forward-looking behavior by smokers; the reason is that data about smoking are readily available, presumably because smoking is widespread and legal, in contrast with other addictive substances, and most aspects of addictive behavior are captured by smoking. Becker et al. (1994) use yearly data from the Tobacco Tax Council for every state and year between 1955-1985 in order to test the main implications of the Becker & Murphy model. By solving a slightly different version of the aforementioned model[48], they derive the following linear difference equation for determining current consumption:

$$C_t = \theta C_{t-1} + \beta\theta C_{t+1} + \theta_1 P_t + \theta_2 e_t + \theta_3 e_{t+1} \text{ (4)},$$

where e_t is a variable measuring the impact of life-cycle events to utility of time t, and the other variables are expressed in standard notation. This equation implies that consumption in different points in time are complements, i.e. that the good is addictive, iff $\theta > 0$. θ measures the degree of reinforcement in consumption, and the higher is θ, the higher is the addictive power of cigarettes. Thus, a clear empirical question arises: whether future consumption is a significant predictor of current consumption, a fact that would suggest forward-looking behavior.

Their empirical strategy consists of performing regressions of per capita cigarette consumption in packs, as derived from state tax-paid sales, on the average price per cigarette pack, the per capita income, variables measuring the smuggling incentives due to tax differentials of adjacent states, and state and time fixed-effects to account for differences in demographic composition between states and for changes in

[48] The model is in discrete, rather in continuous, time, and only the previous period’s consumption, instead of the accumulated capital of past consumption, affects current period’s consumption. However, the main implications of the Becker & Murphy model still hold.

cigarette consumption brought about by new information regarding the health hazards induced by smoking in the 1960's. To correct for endogeneity[49], which arises from the unobservable e_t's, they use past and future prices and excise tax rates as instruments for past and future consumption, since the former are uncorrelated with life-cycle events, but obviously correlated with consumption. Hence, they first run 2SLS regressions using only the past prices to check for the myopic model, and afterwards they include future prices to test for rational addiction.

The results of these regressions unequivocally reject the myopic model of addiction, while providing qualified support for the rational addiction model. The coefficients of past and future consumption are statistically significant and positive, as would be expected if the rational addiction model holds, and the current price coefficient is negative. Furthermore, the long-run elasticity of consumption with respect to a permanent change in price is lower compared to the short-run one[50] (a permanent decrease of 10 percent in price leads to a 4 percent increase in cigarette consumption in the short run, while the increase is almost doubled at 7.5 percent in the long-run), implying that the increase in the short-run will be reinforced through complementarity if $\theta > 0$, leading to a higher long-run response. A similar interpretation should be given to the result comparing the elasticities of anticipated and unanticipated temporary changes in price (consumption increases by 3.5 percent if the change is anticipated, and 3 percent if not, as a response to a 10 percent temporary decrease in price), since if a price change is anticipated, it leads a rational addict to alter his consumption at an earlier date[51]. Many other studies find similar results in support of the rational addiction model, and contrary to the predictions of myopic behavior from smokers (see e.g. Chaloupka (1991)).

However, several implicit assumptions on which the authors rely upon may be subjected to critique, and although they recognized them, they are not addressed in a definitive manner. First of all, a perhaps unrealistic assumption pertaining this analysis is the notion that individuals can forecast future price changes far to the future. While the greatest portion of cigarette price changes can be attributed to

[49] Due to possible serial correlation between them. Even if they were uncorrelated, an error influences the consumption on all dates through the optimizing behavior implied by eq. 4.
[50] Short-run elasticities refer to a permanent change in price at time t that is unknown until then.
[51] Obviously temporary changes have a smaller effect compared to a permanent one, since the latter combines a fall to the current price, as well as to all the future prices.

changes in excise taxes, and these might be expected as a response from the state government to a revenue shortfall, it is highly unlikely that individuals are even able to predict tax changes up to a year in advance, since they are rarely known so much time before enactment, not to mention price changes. Although the authors are aware of the possibility of forecast errors rending their estimates biased, and they run alternative regressions omitting the lead of prices and taxes from the instruments[52], these regressions yield questionable results, even implying negative interest rates or positive coefficients of prices.

In addition, the data are about aggregate cigarette sales, not actual individual consumption, and these two measures are not necessarily the same. For all we know, a rational addict's response to an anticipated future price increase might be to increase purchases, now that the prices are lower, and stockpile them. Although this consideration is redundant if the data are of annual frequency, due to the deteriorating quality of stocked cigarettes, it interacts with the previous one: if the price change is far in the future, no stockpiling occurs, but the change is probably unanticipated; if the price change is in the near term, both anticipation and stockpiling are more probable. In any case, at least one of the two problems remains.

Another field of sensible critique towards the authors regards the matter of correct specification. One such critique is that the regression of the quantity of cigarettes on their price might lead to endogeneity bias. Although, as mentioned above, the greatest part of price changes can be attributed to excise tax changes, the latter cannot account for the whole variation of the former, leaving a part of the variation unexplained and prone to creating an endogeneity bias, since quantity and price are determined simultaneously. A safer strategy, as mentioned by Gruber & Kőszegi (2000), is to use taxes instead of prices, since the latter are truly exogenous regressors. Nevertheless, when they replicate the estimation done by Becker et al. using the same data set, they actually find a positive coefficient on future taxes, contrary to what is expected when adjacent complementarity holds.

In order to fix most of the aforementioned problems and provide a more accurate test for forward-looking behavior, Gruber & Kőszegi (2001) employ a different empirical strategy. They use two datasets for cigarette consumption. The first one is

[52] An additional problem with this strategy is that past prices are poor predictors of future prices.

from the same source used by Becker et al., measuring per capita cigarette consumption derived from state tax-paid sales, although they construct a series of monthly data using state excise tax collections underlying the original time series, which had an annual frequency[53]. However, having collected data on the enactment date of all state cigarette excise tax increases, and of the actual date they went into effect[54], they are able to restrict their attention on consumption changes induced by completely anticipated tax increases. Hence, by examining consumption in the months following the enactment, but preceding the actual effective date, of tax increases, forward-looking behavior can be tested more appropriately, even if individuals are completely unable to predict changes in future prices beyond already announced tax changes.

The second dataset tries to fix the second problem described above, namely that aggregate cigarette sales do not necessarily coincide with actual consumption. They use monthly data from the Vital Statistics Natality database, measuring the per capita number of cigarettes smoked by pregnant mothers, classified by state of residence and month of birth[55]. Although they recognize that this population is by no means representative, since both smoking prevalence and number of cigarettes consumed are lower compared to the typical adult, it is worth examining them, since infant poor health due to maternal smoking is one of the main externalities associated with smoking[56].

The results of the regressions using consumption data from the first dataset on variables measuring the enacted and effective tax rates[57] are as follows: a positive coefficient of the enacted rate, a strongly negative coefficient of current effective rate, and a positive coefficient of the lagged effective rate, with all of them being statistically significant. These findings reveal the following consumption pattern: in anticipation of the future price increase due to the enacted tax rate increase, smokers stockpile cigarettes; when this increase becomes effective, sales fall dramatically at

[53] The time period for which the new series was constructed is from January 1982 to December 1996.
[54] During the period from 1973 to 1996.
[55] The data are from 1989-1996.
[56] Note that even if mothers underreport their actual consumption, this does not lead to systematically biased estimates, unless the underreporting is correlated with tax changes, which seems implausible.
[57] These two measures differ only in the case that a change has been enacted and not yet gone into effect, and these are the cases we are interested in.

the month of the actual increase, due to this stockpiling effect; and the following month, when stocks have been depleted, sales rebound to their previous level. However, this stockpiling effect could be attributed, at least in part, to forward-looking behavior by retailers, and a more definitive test using actual consumption data could be enlightening. When the second dataset is used, only the enacted rate coefficient remains statistically significant and negative, while the others remain negative and statistically insignificant. These evidence imply forward-looking behavior, at least for smoking mothers, rejecting for yet another time the myopic model.

When the authors use alternative specifications, such as fixed trends or differences, the main effects remain the same, albeit somewhat reduced in size. The authors have also used a dummy variable, measuring the presence of a tax change, in the differenced model using natality consumption data. In that way they test whether it is the increase of taxation per se, and not its actual size, that causes pregnant women to reduce cigarette consumption, due to e.g. an increase in antismoking sentiment. However, the dummy is not significant, and the coefficient is not impacted by its inclusion, suggesting that indeed it is the expected increase in taxation that reduces smoking, and that smoking mothers are not myopic.

4.3 The case for time inconsistency

All of the problems pertaining to the Becker et al. (1994) analysis described above are of empirical nature. However, the most important one, that is present in all the relevant literature building upon the Becker & Murphy model, is a theoretical one. More specifically, it is the tacit assumption that agents are time-consistent. As we have already argued, there are many experimental studies indicating that individuals are time-inconsistent, discounting more heavily immediate payoffs compared to ones far to the future.

As we have already implied, there are two ways to distinguish empirically time-consistent from time-inconsistent agents. The first is that the members of the second group search for commitment devices, depending on their degree of sophistication. Knowing that their future self will have different preferences than themselves, agents who recognize their own time-inconsistency, use a wide array of self-control mechanisms to commit themselves to their current decision. The second way is that a time-inconsistent agent has a difficulty to actualize predicted future consumption, depending on their degree of naïvete. Even partially naïve agents are unable to fully appreciate their self-control problems.

There are many studies indicating that smokers very often try to quit believing that the attempt will be successful, only to fail within a very short time period (most of the attempts fail within one week, according to Harris (1993)), or that a considerable proportion of smokers expect to quit in the near future, nevertheless have not yet managed to do so. As Harris (1993) notes, the average smoker tries to quit every eight and a half months, while about 80% of smokers have tried to quit at least once. In the absence of naïvete about future preferences, smoking should have been almost eradicated, even taking into account the different circumstances that every attempt to quit entails.

As another empirical test, specifically for the case of smoking, Gruber & Mullainathan (2001), using data for self-reported well-being from the General Social Survey, find that higher cigarette excise taxes are correlated with higher well-being for smokers, something that does not hold generally for other forms of excise taxation. This finding provides support in favor of time-inconsistency, since time-consistent agents will be made worse off from an increase in taxation, as the model of rational addiction implies. For time-inconsistent smokers, however, taxes are the commitment mechanism they search for, leaving them better off. The desire of smokers, and addicts in general, to enter in rehabilitation programs can also be viewed as evidence of commitment devices employed by time-inconsistent agents.

For all the reasons stated above, Gruber & Kőszegi (2000,2001) model addicts that are quasi-hyperbolic, not exponential, discounters, i.e. they have (β,δ) preferences. Most other features however remain as in Becker & Murphy. The instantaneous utility at time t is additively separable in the consumption of the addictive (α_t) and the non-addictive good: (c_t) $U_t(a_t, c_t, S_t) = v(a_t, S_t) + u(c_t)$. The stock of past consumption of the addictive good evolves as follows: $S_{t+1} = (1-d)(S_t + a_t)$[58], where d is the depreciation rate of the stock. I_t is the income of the agent, p_t the price

[58] This is a small variation from the standard equation used for stock accumulation; however if *d < 1*, these are equivalent through a variable change. Thee departure is done for arithmetic simplicity.

of the addictive good (the non-addictive is treated as the numeraire). Also, v_{as} is assumed to be positive, capturing the notion of addiction (higher past consumption raises future marginal utility).

The authors analyze three types of agents: time-consistent exponential discounters, who maximize standard discounted utility subject to their intertemporal budget constraint and the law of motion of consumption capital; naïve hyperbolic discounters, who perform the same optimization, albeit with a discounted utility as seen in eq. 5, without realizing their self-control problem, i.e. that their future self will overvalue contemporary utility the same way the present self overvalues current utility; and sophisticated hyperbolic discounters, who are aware of that fact. The problem of sophisticates is presented as a game played by the successive selves, and the action space is the vector of consumption (a_t, c_t), where the solution is derived as the subgame perfect equilibrium of that game, as usually presented by the relevant literature.

Using the following quadratic utility functions:

$v(a_t, S_t) = a_a a_t + a_S S_t + \frac{a_{aa}}{2} a_t^2 + \frac{a_{SS}}{2} S_t^2 + a_{as} a_t S_t$ and $u(c_t) = a_c c_t + \frac{a_{cc}}{2} c_t^2$, and assuming that $a_a, a_c, a_{as} > 0$, $a_S, a_{aa}, a_{SS}, a_{SS} < 0$, and that $U_t(a_t, c_t, S_t)$ is strictly concave[59], i.e. that its Hessian is negative definite, it can be proved by backward induction that a_t is linear in S_t for all three types of agents. So $a_t = \lambda_t^i S_t + \mu_t^i$, where λ_t^i and μ_t^i are constants, for *i=TC,n,s* respectively. The authors show that if $\beta \geq \frac{1}{2}$, which is not a far-fetched assumption, then marginal propensities for addiction are stationary for all types of agents, and that $\lambda^{*n} \geq \lambda^{*s} \geq \lambda^{*TC}$, suggesting that a given increase in past consumption raises current consumption most for naïfs, least for time-consistent agents, and somewhere in-between for sophisticates. This means that addiction affects time-consistent agents, who take into account its full effect, less than time-inconsistent ones, and furthermore that naïfs get increasingly worse, relative to sophisticates, at fighting addiction as its level rises. That is because they underestimate the seriousness of the situation they face, believing that they can get 'unhooked' whenever they want, while sophisticates realize how bad their situation is and restrain their impulses.

[59] As noted previously, the fact that $a_{as} > 0$ implies addictive behavior.

The authors show that if the good is sufficiently addictive, an expected increase in future price leads to a decrease in current consumption, since its marginal utility is lowered from the fall in future consumption, for all types of agents. Hence, this model cannot be empirically distinguished from rational models that also suggest forward-looking behavior, despite the significant differences in their underlying assumptions. However, there are radically different policy implications, depending on which model somebody puts his trust on.

4.4 Policy implications

As O'Donoghue & Rabin (2006) argue, analyzing optimal taxation for sin goods, such as unhealthy foods, which resemble addictive ones in the sense that consumption of both bears future health consequences, the optimal tax rate is positive, under fairly general conditions, if agents with even fairly small self-control problems are included in the population, albeit at the slightest of shares. A sin tax provides a self-control device to sophisticates and a sufficient deterrent to naïfs, although it may marginally hurt time-consistent agents. Although this argument seems paternalistic, we cannot rely on the free market to provide such self-control mechanisms, since this would give rise to mechanisms countering self-control demanded by the 'future selves'. Then only private mechanisms remain (e.g. bets), which are not so effective due to enforcement and monitoring reasons.

The somewhat contradictory approach of taxing an individual to increase his welfare was labeled 'libertarian paternalism' by Thaler & Sunstein (2003). However, what distinguishes it from other forms of paternalism, of which economists, not excluding behavioral ones, are wary in general, as the authors view it, is the absence of coercion. Instead, they advocate that in certain cases paternalism is unavoidable, or else we let individuals make decisions that ex-post will regret, or seem inferior according to as objective welfare measures as there exist. As long as the intervention does not restrict their choice set, but simply skews their available choices appropriately, it could well be deemed as palatable even by the most fervent of libertarians.

Returning to the case of smoking, if agents are rational, government intervention is justified only to the extent that it aims to correct externalities induced by smokers to others. However, there is room for significantly higher intervention if the agents are time-inconsistent, because in that case there exist ‘internalities’ that need to be corrected. Calibrating the model presented above, Gruber & Kőszegi (2000,2001) calculate an optimal tax, for the sophisticates case, of at least 1 dollar per pack compared to the average excise tax in the US of 65 cents in 2001. This estimate is quite conservative, since they based their estimate on the fatality costs of smoking, not taking into account health consequences while the smoker is alive, and they imposed parameter values that suggest only modest time-inconsistency. In addition, this recommendation refers only to internality costs associated with smoking, so a tax rate taking externality costs into account should be conceivably higher.

The assumption that individuals in general, and especially smokers, are time-inconsistent has important implications on smoking regulation in various cases. For partially naïve agents the optimal tax is higher compared to fully sophisticate ones, and the higher the degree of naïveté, the higher should be the tax implemented, since for partial naïfs there is not only a self-control problem that needs to be corrected, but also a misrepresentation one, as noted above. If on the other hand smokers are assumed to be time-consistent, then there is very little room for policy intervention, beyond the occurrence of any externalities. Hence, the degree of government intervention justified rises with the degree of naïvete of the consumer, and more naïve agents are in need of more protection compared to more sophisticated ones, who just want an adequate commitment device.

For example, if someone believes that younger smokers are more naïve compared to older smokers, underestimating their quitting probabilities more consistently, then a ban on smoking for non-adults and a tax for older smokers might be a reasonable policy measure. Also, clean air laws might be interpreted not as a measure protecting only the health of non-smokers from second-hand smoke, but also as a device imposing self-control on smokers. Someone could also suggest that a facilitation of admittance to rehabiliation programs and clinics could be an effective measure towards the alleviation of the negative consequences of smoking, or addictive substances in general. The reason is that again they would provide smokers with a much-appreciated self-control device.

Another interesting observation comes from Gruber & Kőszegi (2004), who find that the regressivity of excise cigarette taxes, considered one of their main drawbacks, is lowered by their self-control benefits, especially if low-income individuals have greater self-control problem. Thus, they provide further support to government intervention. While all these results are relevant for most types of addiction, they hold especially true in the case of smoking, where the costs associated with the activity are mainly to the smoker himself, and the harm rises with the level of consumption, unlike e.g. drinking, where there is a threshold level of consumption, after which the negative effects take place.

There exist alternative approaches to addiction modeling, which have not been discussed here. One such approach can be found in Bernheim & Rangel (2004). As the authors argue, there is no empirical factor distinguishing addictive goods compared to any other good in the presence of time-inconsistency, which contradicts many observed behaviors of addicts not found in other situations already analyzed. One such pattern is that addicts usually self-describe their lapses as mistakes[60]; thus, they develop a model containing the possibility of mistakes from the part of the addict[61]. The addict is triggered by certain environmental cues, which lead him to operate in a 'hot' mode, where he does not contemplate all the consequences of his choices. Addiction results from the increased sensitivity of those cues to past experiences[62]. In that case, users of addictive substances may react differently to price changes, depening on whether they are on the 'hot' or 'cold' mode, leading to different welfare implications of a taxation change; in some cases, even a negative tax rate, i.e. a subsidy, is prescribed.

Moreover, it was tacitly assumed that the market is perfectly competitive. An interesting extension would be to introduce market power in those models, since the tobacco industry is highly concentrated. As Becker et al. note, a firm holding market power might price even below its marginal cost, in order that rational addicts be tempted and get 'hooked' to the addictive good, so that the firm will be compensated by the future profits. However, in the case of time-inconsistent agents, it could well be that such a pricing strategy is unnecessary, and that increased advertising alone could

[60] Remember that in many experiments paricipants did not think that they have made a mistake when evidence of their time-inconsistency was presented to them, as discussed in chapter 2.

[61] An almost equivalent assumption is made in the context of a rational addiction model by Orphanides & Zervos (1995), where experimentation with an addictive good leads with some probability to unintentional addiction. An implication is that agents more confident about their susceptibility to addiction might end up addicted. If teenagers were indeed assumed to be particularly overconfident about their ability to resist temptations, this would explain their proclivity for experimentation, and eventual addiction.

[62] This analysis is closely related to Loewenstein's (1996) work on visceral influences.

be enough to get additional consumers ‘hooked’, since naïfs do not recognize the full cost of consuming the good, believing that they will be able to quit if such is their desire. As the preceding analysis has hopefully pointed out by now, the problem of designing appropriate regulatory and tax policies is of particular interest.

Chapter 5

Final discussion: Motivation for bounded rationality

As we have seen in the previous chapter, a model involving boundedly rational agents can be transformed to an equivalent one that assumes perfectly rational agents. This process is known as ‘rationalization’. In the case of addiction, this can be achieved by changing the preferences of agents and transforming them from quasi-hyperbolic to exponential discounters, or by assuming that their past consumption influences their current one, so that addiction emerges through habit formation. The main implication of addiction, forward-looking behavior and fall in current consumption as a reaction to expected price increases, can be equivalently obtained by assuming exponential discounting or habit formation. Spiegler (2011) suggests three ways through which rationalization can be achieved: by assuming that agents are imperfectly informed rather than imperfectly rational, by changing their preferences (as was the case in the addiction example of the previous chapter), or by assuming that there exist certain decision costs that consumers take into account. In each case, we are trying to interpret a choice pattern that seems to originate from a decision error as a rational response to an altered environment.

However, in many cases the non-rational and the ‘rationalized’ models have different welfare implications. An example from Caplin & Schotter (2008) illustrates that point in an excellent way. Imagine an American tourist in London who wants to cross a street. He looks at his left, sees that the road is clear, and when crossing it he gets hit by a car coming from his right. There are two ways to interpret such a result; we could assume that the tourist, without much deliberation, followed a rule-of-thumb that worked in his home environment. Alternatively, we could hypothesize that a certain mental effort is required to perform such a task, and the tourist rationally decides that the decision cost associated with it exceeds the expected benefit of safe crossing. Depending on which assumption one makes, the policy implications differ radically; while a ‘look at your right sign’, like those that actually exist in road crossings in

London, would be beneficial for a boundedly rational tourist, it would not make much of a difference for a rational one that has already incorporated this information into his decision process.

An analogous argument can be made in the case of addiction, studied in the previous chapter: while a high excise tax could serve as the commitment device that time-inconsistent agents eagerly search for, rational addicts would deem this tax as costly in terms of welfare, if it was beyond the internalization of any associated externality. Optimal regulation and policy need to follow different prescriptions as a consequence, depending on the assumption regarding rationality one makes. So a very important question arises: 'Which model should one follow?'; or to put it more blatantly, 'Is there a sufficient reason to abandon the, admittedly highly successful, rationality paradigm?'.

First of all, there is the matter of the descriptive validity of the assumptions regarding rationality. Are individuals perfectly rational? An overwhelming amount of experimental evidence, some of which have been already presented, suggest that they are not, at least not in every situation they face. Individuals discount more heavily outcomes that are closer to the present compared to ones in the distant future, they alter their preferences when presented with inferior alternatives, they change their decisions depending on how the choices are framed, and so forth. Thus, there is a strong reason to believe that 'homo economicus', the rational utility maximizer assumed by standard economic theory, is a concept that fails to accurately depict human decision-making.

However, Friedman (1953) posited that the relevant issue, according to which the appropriateness of the rationality paradigm should be judged, is not whether people are indeed perfectly rational; no one assumes they are. It is rather whether people behave 'as if' they were perfectly rational, so that the theory is evaluated according to the accuracy of its predictions. Or to put it another way, it the predictive success, not the descriptive validity, that should be our criterion for choosing which theory is relevant. We have seen that in many cases, individuals do not behave 'as if' they were perfectly rational: they procrastinate if they do not set deadlines, they overestimate health club attendance, and they embark on diets or quit smoking only temporarily, only to lapse again after a while etc. In all these cases, bounded rationality models are the ones that actually better predict human behavior, not those based on perfect rationality.

A slight modification of the previous argument suggests that people learn to behave optimally through practice, in the end acting as if unboundedly rational. We could therefore avoid this process and model them as perfectly rational from the start. As Conlisk (1996) noted nevertheless, this argument cuts both ways. While learning might work in a simple environment with unchanging circumstances, in the opposite conditions it is actually hindered. He cites the example of life-cycle saving decisions as an example of a particularly complex environment with ever-changing conditions, on which we should not expect especially young individuals to be able to learn effectively. We have already presented examples of people hanging on illiquid assets or procrastinating to transfer their savings to higher interest accounts, which suggest that learning, even if it exists, takes quite a while, and assuming from the start that individuals are rational is not appropriate.

In the end, it seems that bounded rationality in general, and time-inconsistency more specifically, can provide useful insights into human behavior, and we should entertain their use to model it. All the aforementioned reasons explain the existence of a growing body of literature engaging in those issues. However, the bounded rationality territory remains largely unexplored. Ellison (2006) for example, presents a matrix that combining common topics of the Industrial Organization and behavioral biases literatures, and most of those combinations have not yet been analyzed, although they might not all be of particular interest, as he notes. The many cases that bounded rationality and time-inconsistency have been applied to, and their success at providing interesting results, renders us particularly optimistic about the future development of the field.

References

Ainslie, G., Haslam, N., 1992. Hyperbolic Discounting, In Choice over Time, edited by Loewenstein, G. F., Elster, J., 57-92, NY: Russell Sage.

Allais, M., 1953. Le Comportement de l'Homme Rationnel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Americaine. Econometrica 21, 503-546.

Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., Weinberg, S., 2001. The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation. Journal of Economic Perspectives 15, 47-68.

Ariely, D., Wertenbroch, K., 2002. Procrastination, Deadlines, and Performance: Self-Control by Precommitment. Psychological Science 13, 219-224.

Ausubel, L. M., 1991. The Failure of Competition in the Credit Card Market. American Economic Review 81, 50-81

Becker, G.S., 1993. Nobel Lecture: The Economic Way of Looking at Behavior," Journal of Political Economy 101, 385-409.

Becker, G.S., Grossman, M., Murphy, K.M., 1994. An Empirical Analysis of Cigarette Addiction. American Economic Review 84, 396-418.

Becker, G.S., Murphy, K.M., 1988. A Theory of Rational Addiction. Journal of Political Economy 96, 675-700.

Bernheim, B.D., Rangel, A., 2004. Addiction and Cue-Triggered Decision Processes. American Economic Review 94, 1558-1590.

Bohm-Bawerk, E. v. (1889), 1970. Capital and Interest. South Holland: Libertarian Press.

Caplin, A., Schotter, A., 2008. The Foundations of Positive and Normative Economics: A Handbook. Oxford University Press, New York.

Conlisk, J., 1996. Why Bounded Rationality? Journal of Economic Literature 34, 669-700

Chaloupka, F., 1991. Rational Addictive Behavior and Cigarette Smoking. Journal of Political Economy 99, 722-742.

DellaVigna, S., 2009. Psychology and Economics: Evidence from the Field. Journal of Economic Literature 47, 315-372.

DellaVigna, S., Malmendier, U., 2004. Contract Design and Self-Control: Theory and Evidence. Quarterly Journal of Economics 119, 353-402.

DellaVigna, S., Malmendier, U., 2006. Paying Not to Go to the Gym. American Economic Review 96, 694-719.

Donegan, N.H., Rodin, J., O'Brien, C.P., Solomon, R.L., 1983. A Learning Theory Approach to Commonalities, In Commonalities in Substance Abuse and Habitual Behavior, edited by Levison, P. K., Gerstein, D. R., Maloff D. R., Lexington, Mass.: Heath

Duesenberry, J., 1949. Income, Saving, and the Theory of Consumer Behavior. Cambridge, MA: Harvard University Press.

Eliaz, K., Spiegler, R., 2006. Contracting with Diversely Naive Agents. Review of Economic Studies 73, 689-714.

Eliaz, K., Spiegler, R., 2011. Consideration Sets and Competitive Marketing. Review of Economic Studies 78, 235-262.

Ellison, G., 2006. Bounded Rationality in Industrial Organization, In Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, edited by Blundell, R., Newey, W. K., and Persson, T., Vol II, 142-174, Cambridge University Press, New York.

Frederick, S., Loewenstein, G. F., O'Donoghue, T., 2002. Time Discounting and Time Preference: A Critical Review. Journal of Economic Literature 40, 351-401.

Friedman, M., 1953. The Methodology of Positive Economics, in The Philosophy of Economics: An Anthology 2, 180-213

Gruber, J., Kőszegi, B., 2000. Is Addiction ‘Rational’? Theory and Evidence. NBER Working Paper 7507.

Gruber, J., Kőszegi, B., 2001. Is Addiction ‘Rational’? Theory and Evidence. Quarterly Journal of Economics 116, 1261-1305.

Gruber, J., Kőszegi, B., 2004. Tax Incidence when Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes. Journal of Public Economics 88, 1959-1987.

Gruber, J., Mullainathan, S., 2005. Do Cigarette Taxes Make Smokers Happier? Advances in Economic Analysis & Policy Volume 5.

Gul, F., Pesendorfer, W., 2001. Temptation and Self-Control, Econometrica 69, 1403-1435.

Harris, J.E., Deadly Choices: Coping with Health Risks in Everyday Life. New York: Basic Books, 1993.

Hoch, S. J., Loewenstein, G. F., 1991. Time-Inconsistent Preferences and Consumer Self-Control. Journal of Consumer Research 17, 492-507.

Kahneman, D., Knetsch, J. L., Thaler, R. H., 1990. Experimental Tests of the Endowment Effect and the Coase Theorem. Journal of Political Economy 98, 1325-1348.

Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. Econometrica 47, 263-291.

Koopmans, T. C., 1960. Stationary Ordinal Utility and Impatience. Econometrica 28, 287-309.

Kőszegi, B., Rabin, M., 2006. A Model of Reference-Dependent Preferences. Quarterly Journal of Economics 121, 1133-1165.

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

Loewenstein, G. F., 1987. Anticipation and the Valuation of Delayed Consumption. Economic Journal 97, 666-684.

Loewenstein, G. F., 1996. Out of Control: Visceral Influences on Behavior. Organizational Behavior and Human Decision Processes 65, 272-292.

Loewenstein, G. F., O'Donoghue, T., Rabin, M., 2003. Projection Bias in Predicting Future Utility. Quarterly Journal of Economics 118, 1209-1248.

Loewenstein, G. F., Prelec, D., 1992. Anomalies in Intertemporal Choice: Evidence and an Interpretation. Quarterly Journal of Economics 107, 573-597.

Loewenstein, G. F., Prelec, D., 1993. Preferences for Sequences of Outcomes. Psychological Review 100, 91-108.

Loewenstein, G. F., Sicherman, N., 1991. Do Workers Prefer Increasing Wage Profiles? Journal of Labor Economics 9, 67-84.

Mullainathan, S., Schwartzstein, J., Shleifer, A., 2008. Coarse Thinking and Persuasion. Quarterly Journal of Economics 123, 557-619.

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

O'Donoghue, T., Rabin, M., 2001. Choice and Procrastination. Quarterly Journal of Economics 116, 121-160.

O'Donoghue, T., Rabin, M., 2006. Optimal Sin Taxes. Journal of Public Economics 90, 1825-1849.

Orhanides, A., Zervos, D., 1995. Rational Addiction with Learning and Regret. Journal of Political Economy 103, 739-758.

Phelps, E. S., Pollak, R. A., 1968. On Second-Best National Saving and Game-Equilibrium Growth. Review of Economic Studies 35, 185-99.

Pollak, R. A., 1968. Consistent Planning. Review of Economic Studies 35, 201-208.

Pollak, R. A., 1970. Habit Formation and Dynamic Demand Functions. Journal of Political Economy 78, 745-763.

Rae, J., 1834. The Sociological Theory of Capital. London: Macmillan.

Read, D., 2001. Is Time Discounting Hyperbolic or Subadditive? Journal of Risk and Uncertainty 23, 5-32.

Read, D., Loewenstein, G. F., Rabin, M., 1999. Choice Bracketing. Journal of Risk and Uncertainty 19, 171-197.

Rubinstein, A., 2003. "Economics and Psychology"? The Case of Hyperbolic Discounting. International Economic Review 44, 1207-1216.

Rubinstein, A., 1991. Modeling Bounded Rationality. MIT Press.

Samuelson, P., 1937. A Note on Measurement of Utility. Review of Economic Studies 4, 155-161.

Samuelson, W., Zeckhauser, R., 1988. Status Quo Bias in Decision Making. Journal of Risk and Uncertainty 1, 7-59.

Simon, H. A., 1955. A Behavioral Model of Rational Choice. Quarterly Journal of Economics 59, 99-118.

Simon, H. A., 1956. Rational Choice and the Structure of the Environment. Psychological Review 63, 129-138.

Spiegler, R., 2006. The Market for Quacks. Review of Economic Studies 73, 1113-1131.

Spiegler, R., 2011. Bounded Rationality and Industrial Organization. Oxford University Press, New York.

Strotz, R. H., 1955. Myopia and Inconsistency in Dynamic Utility Maximization. Review of Economic Studies 23, 165-180.

Thaler, R. H., 1981. Some Empirical Evidence on Dynamic Inconsistency. Economics Letters 8, 201-207.

Thaler, R. H., 1990. Anomalies: Saving, Fungibility, and Mental Accounts. Journal of Economic Perspectives 4, 193-205.

Thaler, R. H., Sunstein, C. R., 2003. Libertarian Paternalism. American Economic Review 93, 175-179

Tversky, A., 1977. Features of Similarity. Psychological Review 84, 327-352.

Tversky, A., Kahneman, D., 1981. The Framing of Decisions and the Psychology of Choice. Science 211, 453-458.