Present-biased preferences, retirement planning and demand for commitments

Mario Bernasconi*

October 2020

Preliminary version. Do not quote or circulate.

Abstract

Evidence of present-biased preferences is spread over different fields of economics. However, most of the literature concerning life-cycle models and present biased preferences abstracted from retirement planning. Building on Diamond and Koszegi (2003), this paper considers a three-period stylized model where present-biased agents decide how much to consume and when to retire. Agents unaware of their bias exhibit time-inconsistent choices, which can result in early retirement with respect to their original plan. Agents who are aware of their bias, instead, can over-save or under-save to influence their future retirement behaviour. These agents can also fully commit themselves to their optimal consumption plan using an illiquid asset as a commitment device when the retirement date is exogenously set. However, when trying to address consumption and retirement choices together, illiquid assets are not always a perfect commitment device. Partially naïve agents could even demand too much of the commitment and reduce their welfare.

1 Introduction

People are impatient in the sense that they value 10$ today more than 10$ tomorrow. Traditionally, economic theory modelled this aspect of agents' preferences assuming that people use an exponential function to discount future utility pay-offs. Exponential discount functions induce time consistent preferences: an agent choice in period 0 regarding time $\tau > 1$ does not change if the same decision regarding τ is taken in period $\tau - 1$, given that we consider an uncertainty-free context (choice reconsiderations are not due to revealed uncertainty). A rational agent with time consistent preference decides his optimal consumption and labour supply path in the starting period $t = 0$ for all future periods $t = 0, ..., T$ and will always stick to his initial decision because in every period t the maximization problem would give the same solution (consistency). In such models, the attribute of self-control does not play any role.

However, a variety of experimental evidence suggests that people have time-inconsistent preferences (e.g. Thaler (1981)). This means that people are not just impatient, but they also tend to have present-biased preferences. In O'Donoghue and Rabin (1999)'s words, when considering trade-offs between two future moments, present-biased preferences give stronger relative weight to the earlier moment as it gets closer. This explains why people tend to procrastinate immediate-cost activities and preproperate immediate-reward activities.

These present-biased and time-inconsistent preferences have been modelled in the literature with hyperbolic discounting functions. Hyperbolic discounting exactly captures the fact that

*Department of Econometrics and Operations Research, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands (e-mail: m.bernasconi@tilburguniversity.edu)

a person has a declining rate of time preference. Since hyperbolic discount functions induce dynamically inconsistent preferences, they imply a motive for consumers to constrain their own future choices. This is something that a rational agent with time-consistent preferences would never do. For analytical convenience, time inconsistent preference are usually modelled with quasi-hyperbolic discounting functions rather than more psychologically accurate hyperbolic discounting functions. For example, Laibson (1997) influential paper model these preferences and emphasizes the use of illiquid assets as a commitment technology. In his model, illiquid assets can be used to constrain agents' future consumption and savings and to limit time inconsistent choices. However, this model assumes that labour is supplied inelastically.

Surprisingly, little attention has been devoted to labour supply choices of time inconsistent agents. There exist only few studies where labour supply is assumed to be endogenous and they are not conclusive. Diamond and Koszegi (2003) propose a basic three-periods model for the analysis of the effect of endogenous retirement decisions on savings behaviour in a quasi-hyperbolic discounting context. However, in a follow-up to Diamond and Koszegi (2003), Holmes (2010) proved numerically that in their model retirement plans are not time-inconsistent. More recently, Feigenbaum and Findley (2013) and Feigenbaum and Findley (2015) further investigate how present-biased people make retirement choices and question Holmes (2010) conclusion. Nevertheless, they only analysed the behaviour of naïve agents and not of sophisticated agents.

Present-biased preferences provide a motive for commitments to affect both consumption and retirement choices. I investigate the effect of a commitment device in a setting where more than one decision has to be taken. I follow Laibson (1997) and consider illiquid assets as a possible commitment device. In my model, illiquid assets are a perfect commitment device when they are used to influence the agent's own behaviour only with respect to consumption. However, I show that when two interdependent choices have to be taken, such as consumption and retirement, the demand for commitment is non-trivial. In some cases, a rational agent who is aware of his time-inconsistent preferences is able to use the illiquid asset to commit both of his (future) actions. In other cases, the commitment might push one action in an unintended way while trying to address the other. In such cases we might expect individuals to demand less commitments. However, a rational agent could either prefer to limit his use of the illiquid asset to avoid the unintended consequences, or he could rely on the commitment even more to counterbalance its negative consequences.

Evidence for self-commitment is spread over different fields but it is mixed. For example, Giné et al. (2010) found that only a minority of smokers are willing to take up the voluntary commitment product the authors designed. Ashraf et al. (2006) and Beshears et al. (2015) provide evidence from field experiments that people do use illiquid saving accounts as a commitment device to increase savings. Kaur et al. (2010) finds evidence of self-control problems and commitment effectiveness in the workplace context. It is also not clear whether people are willing to pay for commitment devices. Augenblick et al. (2015) find that commitment is popular in their student sample at a zero price but not at a strictly positive price. Laibson (2015) theoretical model originates from the observation that few people express a willingness to pay a significant price to have their choice-set reduced. In his model, a small price of commitment can tip the scales against commitment. On the other hand, Schilbach (2019) provides evidence from a field experiment in India that participants exhibited significant demand for commitment to sobriety, even at the cost of giving up considerable payments. The author also found convincing evidence that the reduction of day-drinking increased participants savings. This suggests that with one commitment device the participants to the experiment were able to achieve two interdependent results and thus the marginal utility deriving from committing was particularly high implying a reason to pay a positive price.

Motivated by these findings, I suggest that in order to better understand the empirical evidence about the demand for commitments we should consider that often individuals take decisions that are interdependent. If one commitment device is able to address multiple actions in a favourable way at the same time it is more appealing. On the other hand, if the commitment pushes one decision in an unindented direction while trying to address another decision it is less appealing.

My first contribution originates from the lack of a theoretical work that endogenises labour supply and investigates the demand for illiquid assets of sophisticated agents. A second contribution concerns the literature regarding self-commitment devices. I suggest that part of the reason for the mixed evidence about demand for commitments is due to the fact that researches usually assume that individuals are only trying to address one particular decision at a time, such as quitting smoking or drinking, saving, going to the gym.

The remainder of the paper is organized as follows: in section 2 the model is presented, in section 3 and 4 I solve the model respectively for the cases of exogenous and endogenous labour supply, in section 5 the possibility of buying an illiquid asset is introduced, in section 6 the conclusions are drawn.

2 Model

Time Preferences

Samuelson (1937) posed the starting point of how economists model inter-temporal choices: the discounted utility model. The model prescribes that an agent's inter-temporal utility can be described as

$$U(c_0, ..., c_T) = \sum_{t=0}^{T} D(t)u(c_t) \tag{1}$$

$$\text{with } D(t) = \left(\frac{1}{1+\rho}\right)^t \tag{2}$$

where $u(c_t)$ is the instantaneous utility deriving from consumption in period t and ρ is the agent's rate of time preference or discount rate. The model is particularly attractive because all the conflicting motives that concern inter-temporal choices are captured by a single parameter. The common formulation is then $D(t) = \beta^t$ with $\beta \in (0, 1)$. This functional form implies a constant discount factor between any two equally spaced periods:

$$\frac{D(0)}{D(1)} = \frac{1}{\beta} = \frac{D(t)}{D(t+1)} \quad \forall t \tag{3}$$

There are, however, many phenomena that cannot be explained by the classic discounted utility model. One of the most documented of these phenomena is hyperbolic discounting. Hyperbolic discounting means that time preferences exhibit a declining rate. A widely used example to present this feature is that of preference reversal: people that prefer \$110 in 31 days over \$100 in 30 days, but also prefer \$100 now over \$110 tomorrow (see Frederick et al. (2002) for an exhaustive literature review on the topic). These choices imply that the discount factor between two equally spaced periods is not constant.

Motivated by these empirical findings, Phelps and Pollak (1968) proposed a discounting function that can accommodate the qualitative features of hyperbolic discounting:

$$D(t) = \begin{cases} 1 & \text{if } t = 0 \\ \beta\delta^t & \text{if } t > 0 \end{cases} \quad \text{with } \beta, \delta \in (0, 1) \tag{4}$$

$$\frac{D(0)}{D(1)} = \frac{1}{\beta\delta} > \frac{1}{\delta} = \frac{D(t)}{D(t+1)} \quad \forall t > 0 \tag{5}$$

(5) shows how this functional form implies a higher discount rate between today and tomorrow, but a lower constant discount rate for all the other subsequent periods.

Before proceeding with the model, let us define via an example what we mean by naïveté and sophistication and by time-inconsistency. Consider the following game: an agent at time $t = 0$ (today) has to decide whether he wants to receive 100$ today (option A) or 110$ tomorrow (t=1) (option B). Also, at time $t = 0$ he has to decide whether he wants to receive 100$ in $t = 30$ (option C) or 110$ in $t = 31$ (option D). The latter choice is not binding. When the time comes, at $t = 30$ he will be able to change his previous decision. The choice of A and D could not happen in a model where we assume that the agent discounts future utility exponentially:

$$A \succ B \text{ i.e. } u(100) > \beta u(110) \tag{6}$$

$$C \prec D \text{ i.e. } \beta^{30} u(100) < \beta^{31} u(110) \tag{7}$$

$$\implies \beta u(110) < u(100) < \beta u(110) \;\lightning \tag{8}$$

Suppose that the present-biased agent at $t = 0$ chooses A and D, which means:

$$A \succ B \text{ i.e. } u(100) > \beta\delta u(110) \tag{9}$$

$$C \prec D \text{ i.e. } \beta\delta^{30} u(100) < \beta\delta^{31} u(110) \tag{10}$$

$$\implies \beta\delta u(110) < u(100) < \delta u(110) \tag{11}$$

The very same agent, when $t = 30$ comes, given his preferences, will change his decision from D to C. This agent has time-inconsistent preferences and time-inconsistent behaviour. An agent is naïve when he is not aware of the structure of his time preferences: he does not realize that at $t = 30$ he will change his mind. Given that at time $t = 0$ he discounts pay-offs only by a factor δ from $t = 31$ to $t = 30$, he thinks that he will do the same when time $t = 30$ comes.

On the other hand, a sophisticated agent would still prefer D to C at time $t = 0$, but already knows that at time $t = 30$ he would prefer C. Therefore he realizes that D is ultimately not an option and he would immediately pick C at time $t = 0$. Therefore, he cannot, by definition, have time-inconsistent behaviour, even if his time-preferences are not consistent.

Definition 1. *In an uncertainty-free context, an agent has time-inconsistent preferences when his preferences over a choice to be taken in period* 0 *regarding time* $\tau > 1$ *are different from his preferences if the same decision regarding* τ *is taken in period* $\tau - 1$.

Definition 2. *In an uncertainty-free context, an agent has time-inconsistent behaviour when his expectation regarding his future behaviour are incorrect and so his planning is not optimal.*

Definition 3. *An agent is sophisticated (naïve) if he is (not) aware of the structure of his time-preferences.*

Sophistication can be a desirable attribute. In particular, time-inconsistent behaviour can lead to additional costs for the agent that ex-post he would have preferred to avoid. A classic example would be gym membership and attendance: people seem to overestimate their future gym attendance and thus the choice to buy a flat rate contract (monthly or annual subscriptions) leads to high costs with low benefits (DellaVigna and Malmendier, 2006). Even more important are the decisions regarding savings and retirement, which involve forward looking behaviour and can greatly affect people's welfare. In that case, time-inconsistent planning can entail large transaction costs or opportunity costs. For example, people might regret not working enough when they were young.

Model Setup

The model is built on that of Diamond and Koszegi (2003), which was adapted by the authors to introduce the retirement decision to that of Laibson (1997) and also to make it more tractable. This setup is the one analysed also in Holmes (2010) and Feigenbaum and Findley (2013).

The consumer is characterized by a constant relative risk aversion (CRRA) instantaneous utility function. The inter-temporal choices are affected by a quasi-hyperbolic discounting function characterized by two parameters (β, δ). In order to simplify the analysis, I assume that the consumer has a unitary relative risk aversion parameter, i.e.:

$$u(c) = \ln(c) \tag{12}$$

The theoretical implication of the first assumption (CRRA) is that the consumer's decisions are unaffected by scale. That is, the fraction of wealth optimally consumed in one period is independent of the level of initial wealth. Consider a two-periods allocation problem where the consumer has to decide how much to consume in the two periods given a discounted total wealth of W and return $R = 1 + r$ which is constant and exogenous:

$$\max_{c_1, c_2} u(c_1) + \beta\delta u(c_2) \text{ s.t. } c_1 + \frac{1}{R}c_2 = W \tag{13}$$

$$\max_{c_1} u(c_1) + \beta\delta u(R(W - c_1)) \tag{14}$$

$$\frac{\partial U}{\partial c_1} = \frac{1}{c_1} - R\beta\delta\frac{1}{R(W - c_1)} = 0 \tag{15}$$

$$c_1 = \frac{1}{1 + \beta\delta}W = \lambda(\beta, \delta)W \tag{16}$$

$$c_2 = \frac{\beta\delta}{1 + \beta\delta}WR = (1 - \lambda(\beta, \delta))WR \tag{17}$$

The result is that consumption in the first period is always a fixed share of wealth, with $0 < \lambda < 1$. In particular, since we assume $\beta, \delta \in (0, 1)$, it follows that $\lambda > 1 - \lambda$. This property will be used many times later in the paper omitting the derivation. The CRRA assumption is made to simplify the exposition of the model. The same is true for assuming a quasi-hyperbolic discounting function instead of psychologically more accurate hyperbolic discounting.

I analyze a three-periods model. In period 1 the agent starts with an initial endowment $w_0 > 0$, he has to work and earn a salary w_1. He can consume c_1 and saves $s_1 = w_0 + w_1 - c_1$. In period 2 he can decide to work or to retire. If he works he earns a salary w_2 and faces an additive disutility due to his effort e; if he doesn't work he has at his disposal only the savings s_1R. He consumes c_2 and decides how much to leave for the next period. In the third period he cannot work and he clearly consumes everything that is left from the previous periods.

In a context with time consistent preferences (such as exponential discounting) the consumer would choose an optimal consumption path (c_1, c_2, c_3) in period $t = 1$ and would later always stick to that path because in every later period he would find it optimal. Things are different with time inconsistent preferences. If the consumer is naïve, that is he is not aware of his time inconsistent preferences, then he would choose an optimal consumption path in the initial period but he would not stick to it in the later periods. If the consumer is sophisticated, i.e. he is aware of his time inconsistent preferences, he anticipates his future inconsistency and can try to affect his future behaviour. In particular, if commitment devices are available, he would use these devices to constrain his future behaviour.

3 Consumption Decision

The consumer's choice can be modelled as an equilibrium in a sequential game played by the different selves. The game can then be solved using backward induction starting from the last period. The agent in the model understands perfectly the consequences of his actions, and acts optimally within the constraints imposed by his discount function. There is no uncertainty in this setting. Let's start by ignoring the retirement decision to show time inconsistency regarding the consumption decision. The consumption decision has already been studied extensively and my work does not provide any particular novelty in that sense, apart from studying it in a stylized model. The only new result is that with logarithmic utility the consumption path of a naïve agent is observationally equivalent to that of a sophisticated agent if labour is supplied exogenously and no commitments are available. This property makes the analysis of the behaviour of sophisticated agent who can decide when to retire easier.

Let's consider the consumption decision in a stylized version of the model of Laibson (1997). Consider a setting where the agent only works in period 1, gets a wage w_1 and the only choice he can make is about consumption.

Naïve Agent

Consider a naïve agent that in period $t = 1$ has to decide his optimal consumption path by solving:

$$\max_{c_1,c_2,c_3} u(c_1) + \beta\delta u(c_2) + \beta\delta^2 u(c_3) \text{ s.t. } c_1 + \frac{1}{R}c_2 + \frac{1}{R^2}c_3 = w_0 + w_1 \tag{18}$$

He knows that in the last period he would just consume what is left. In the second period he would consume a fixed share of wealth, as I showed earlier, and he leaves the rest for the third period. Since he is naïve, he thinks that in period 2 he will have the same time preference as he has now (in period 1). Therefore he concludes that in period 2 he would decide based on the following maximization problem:

$$\max_{c_2,c_3} \beta\delta u(c_2) + \beta\delta^2 u(c_3) = \max_{c_2,c_3} u(c_2) + \delta u(c_3) \text{ s.t. } c_2 + \frac{1}{R}c_3 = s_1 R \tag{19}$$

$$\implies c_2^* = \frac{1}{1+\delta}s_1 R,\ c_3^* = \frac{\delta}{1+\delta}s_1 R^2 \tag{20}$$

with s_1 being the saving from period 1: $s_1 = w_0 + w_1 - c_1$. The choice of c_1/s_1 is made in period 1 by solving the following:

$$\max_{s_1} u(w_0 + w_1 - s_1) + \beta\delta u\left(\frac{1}{1+\delta}s_1 R\right) + \beta\delta^2 u\left(\frac{\delta}{1+\delta}s_1 R^2\right) \tag{21}$$

$$\implies c_1^* = \frac{1}{\beta\delta^2 + \beta\delta + 1}(w_0 + w_1)\ \&\ s_1^* = \frac{\beta\delta^2 + \beta\delta}{\beta\delta^2 + \beta\delta + 1}(w_0 + w_1) \tag{22}$$

When self 1 is naïve his optimal consumption plan is (c_1^*, c_2^*, c_3^*). However, because of quasi-hyperbolic discounting, he will not stick to this plan. Self 1 will indeed consume c_1^* and save what is left, so that self 2 will inherit a wealth $s_1^* R$. However, self 2 will decide how much to consume based on the following:

$$\max_{c_2,c_3} u(c_2) + \beta\delta u(c_3) \text{ s.t. } c_2 + \frac{1}{R}c_3 = s_1 R \tag{23}$$

$$\implies c_2^{**} = \frac{1}{1+\beta\delta}s_1 R,\ c_3^{**} = \frac{\beta\delta}{1+\beta\delta}s_1 R^2 \tag{24}$$

It is easy to see that, since $\beta, \delta \in (0, 1)$, we have $c_2^* < c_2^{**}$ and $c_3^* > c_3^{**}$. This implies that if self 1 is naïve he would plan to consume a certain quantity in period 2 but would end up consuming more because of his present-time bias. The opposite is true for period 3. By definition, this result in time-inconsistent behaviour.

Result 1. *Quasi-hyperbolic discounting induces time-inconsistent consumption planning of a naïve agent.*

Sophisticated Agent

On the other hand, a sophisticated self 1 knows that self 2 would ultimately consume (c_2^{**}, c_3^{**}) regardless of his initial planning. His choice is then the solution of the problem in (25). With instantaneous logarithmic utility, problem (21) and problem (25) give the same solution. This means that a naïve self 1 would make the same choice of a sophisticated self 1. However, this does not mean that a sophisticated agent is time-inconsistent. In fact, his planned consumption path is equal to the realized one, because already in period 1 he knows that in periods 2 and 3 he will consume the optimal quantity chosen by self 2, which is different from self 1's best option.

$$\max_{s_1} u(w_0 + w_1 - s_1) + \beta\delta u\left(\frac{1}{1+\beta\delta} s_1 R\right) + \beta\delta^2 u\left(\frac{\beta\delta}{1+\beta\delta} s_1 R^2\right) \tag{25}$$

$$\implies c_1^* = \frac{1}{\beta\delta^2 + \beta\delta + 1}(w_0 + w_1) \ \& \ s_1^* = \frac{\beta\delta^2 + \beta\delta}{\beta\delta^2 + \beta\delta + 1}(w_0 + w_1) \tag{26}$$

Result 2. *With logarithmic utility and exogenous labour supply, the consumption path of a naïve agent is observationally equivalent to that of a sophisticated agent.*

This result highlights that it might not be possible to infer sophistication from observed choices when labour is supplied inelastically. This property of logarithmic utilities is useful to simplify the analysis of the model when retirement is made endogenous. In particular, this results highlight that the optimal consumption and saving level of self 1 for period 1 is the same regardless of the consumption allocation between periods 2 and 3.

4 Retirement Decision

The retirement decision has been first added to a model with quasi-hyperbolic discounting by Diamond and Koszegi (2003). Their focus was primarily on comparative statics in a stylized three-period model and on observational equivalence between quasi-hyperbolic and exponential discounting. They showed that for certain levels of savings an agent could exhibit time-inconsistency regarding retirement choices, but they did not formally verify whether such levels of savings could arise in the model. Holmes (2010) commented on their work in a follow-up paper and showed the "non-existence of inconsistent retirement planners" via a numerical example. However, as Holmes (2010) himself reported, he did not solve a general version of the model. He imposed some simplifying assumptions that can be summarized as follow in terms of my notation: $R = 1$, $\delta = 1$, $w_0 = 0$, $w_1 = w_2 = 1$. In particular, assuming a flat path for labour income and zero initial endowment lead to the result that time-inconsistency is not possible. This is because the resulting savings will not be high enough to induce early retirement. Zero initial endowment was also assumed in Diamond and Koszegi (2003), but it is actually necessary to have either a high enough positive initial endowment or a higher salary in the first period for the retirement plan to be inconsistent. I will solve the model without imposing particular restrictions and I will fully characterize the behaviours and expectations of naïve and sophisticated agents.

Naïve Agent

Let's now endogenize the retirement decision in period 2. Self 2 inherits $s_1 R$, the savings of self 1 times the (constant) return, which he takes as given. He can decide to work and earn w_2 and incurs an additive constant disutility cost given by the effort $e > 0$, or he can retire. The agent is naïve and not aware of his time inconsistent preferences. If self 2 does not work the wealth he can spend is just $s_1 R$ and his utility is given by:

$$u(\lambda_2 s_1 R) + \beta\delta u((1 - \lambda_2) s_1 R^2) \tag{27}$$

If he does work, instead, he gets:

$$u(\lambda_2(s_1R + w_2)) - e + \beta\delta u((1-\lambda_2)(s_1R + w_2)R) \tag{28}$$

Self 2 then works if:

$$u(\lambda_2(s_1R + w_2)) - u(\lambda_2 s_1 R) + \beta\delta[u((1-\lambda_2)(s_1R + w_2)R) - u((1-\lambda_2)s_1R^2)] \geq e \tag{29}$$

Here $\lambda_2 = \frac{1}{1+\beta\delta}$ because self 2 is deciding on the allocation. Since $u(\cdot)$ is a concave function the condition expressed by (29) holds for every s_1 smaller than a certain threshold. Let's define this threshold as $\bar{k}_2$, that is the value of s_1 for which self 2 is indifferent between working or not (left side of (29) equal to the right side). Self 2 works if $s_1 \leq \bar{k}_2$. This is intuitive: if my savings are small I have to work when I am old, whereas if I saved enough I can retire earlier.

$$\bar{k}_2 = \frac{w_2}{R[\exp(\frac{e}{1+\beta\delta}) - 1]} \tag{30}$$

This shows how self 2 is more likely to work for large values of wage (w_2) and of the discount function (large β, δ, that is he cares more about the future). The opposite is true if the interest rate is large (R) and the disutility from working is high (e).
Let's now consider self 1 point of view regarding the decision in period 2. Since self 1 is naïve, he thinks that self 2 will stick to the allocation that is optimal for self 1 ($\lambda_1 = \frac{1}{1+\delta}$) and that self 2 will decide based on the following comparison:

$$\beta\delta[u(\lambda_1(s_1R + w_2)) - u(\lambda_1 s_1 R)] + \beta\delta^2[u((1-\lambda_1)(s_1R + w_2)R) - u((1-\lambda_1)s_1R^2)] \geq \beta\delta e \tag{31}$$

which can be rewritten as:

$$u(\lambda_1(s_1R + w_2)) - u(\lambda_1 s_1 R) + \delta[u((1-\lambda_1)(s_1R + w_2)R) - u((1-\lambda_1)s_1R^2)] \geq e \tag{32}$$

The difference between (29) and (32) is just a factor $\beta \in (0,1)$, which generates time-inconsistent preferences, and the different λ's are actually innocuous because they cancel out with logarithmic utility. The left-hand side of (32) is always greater than the left-hand side of (29). Consequently, there is a range of values of s_1 for which self 2 would not work, but self 1 would like him to. Naïve self 1 thinks that the threshold that self 2 will use to decide is given by:

$$\bar{k}_1 = \frac{w_2}{R[\exp(\frac{e}{1+\delta}) - 1]} \tag{33}$$

$\beta \in (0,1)$ implies that $\bar{k}_1 > \bar{k}_2$. In particular, for $\bar{k}_2 < s_1 \leq \bar{k}_1$ self 2 will retire, but naïve self 1 thinks self 2 will work.

Possibility of Time-Inconsistent Behaviour

The previous section showed how high savings could induce early retirement. In particular, in order to have time inconsistency the savings should be low enough so that self 1 thinks that self 2 will work, but high enough so that self 2 prefers not to work. This level of savings can in fact originate from the choice of self 1. We would need $s_1^{*w} \in (\bar{k}_2, \bar{k}_1]$, which ultimately depends on the relationship between the initial wealth and the salaries:

$$\frac{w_2}{R[\exp(\frac{e}{1+\beta\delta}) - 1]} < \frac{\beta\delta^2(w_0 + w_1) + \beta\delta(w_0 + w_1) - w_2/R}{\beta\delta^2 + \beta\delta + 1} \leq \frac{w_2}{R[\exp(\frac{e}{1+\delta}) - 1]} \tag{34}$$

$$\frac{1}{(\beta\delta + \beta\delta^2)R}\left[\frac{(\beta\delta^2 + \beta\delta + 1)}{\exp(\frac{e}{1+\beta\delta}) - 1} + 1\right] < \frac{w_0 + w_1}{w_2} \leq \frac{1}{(\beta\delta + \beta\delta^2)R}\left[\frac{(\beta\delta^2 + \beta\delta + 1)}{\exp(\frac{e}{1+\delta}) - 1} + 1\right] \tag{35}$$

Given $\beta, \delta \in (0,1)$ the LHS of (35) is indeed lower than its RHS and both can attain positive and negative values over $\mathbb{R}$. Therefore, there exists a value $\frac{w_0+w_1}{w_2} \in \mathbb{R}^+$ that satisfy the condition.

Result 3. *Quasi-hyperbolic discounting can induce time-inconsistent behaviour regarding both consumption and retirement and can induce earlier retirement for a naïve agent.*

Sophisticated Agent

So far, I showed how quasi-hyperbolic discounting induces time-inconsistent behaviour regarding both consumption and the retirement decision of a naïve agent. By definition, a sophisticated agent cannot have time-inconsistent behaviour: he knows that his preferences are present-biased and his expectations are aligned with his actual behaviour. However, quasi-hyperbolic discounting will still affect self 1 behaviour if he is sophisticated because there is a conflict between self 1's optimal choice and self 2's optimal choice and self 1 might try to affect the decision of self 2. The following discussion will explain this possibility.

Self 1 is sophisticated and knows that self 2 will work if $s_1 \leq \bar{k}_2$ and will retire otherwise. This means that self 1 can influence the time of retirement by choosing the level of savings that self 2 will inherit. Using backward induction self 1 knows how self 2 will behave and thus he maximizes his discounted lifetime utility only with respect to s_1 (or c_1, that is equivalent). We can formalize this as:

$$\max_{s_1} U(s_1) \text{ s.t. budget constraint} \tag{36}$$

where $U(s_1)$ is the lifetime utility from periods 1, 2 and 3. For $s_1 \leq \bar{k}_2$, $U^w(s_1)$ would be the lifetime utility deriving from working in period 2. Vice-versa, for $s_1 > \bar{k}_2$, $U^r(s_1)$ would correspond to retirement in period 2. Suppose now that self 2 is forced to work in period 2. Then self 1 would choose the optimal savings level s_1^{*w}. If self 2 is forced to retire in period 2, instead, self 1 would choose s_1^{*r}. It is intuitive that $s_1^{*r} > s_1^{*w}$, because an agent who is forced to retire earlier saves more than one that is forced to retire later. In particular the savings have the following functional form:

$$s_1^{*w} = \frac{\beta\delta^2(w_0 + w_1) + \beta\delta(w_0 + w_1) - w_2/R}{\beta\delta^2 + \beta\delta + 1} \tag{37}$$

$$s_1^{*r} = \frac{\beta\delta^2(w_0 + w_1) + \beta\delta(w_0 + w_1)}{\beta\delta^2 + \beta\delta + 1} \tag{38}$$

Imagine now that self 1 can somehow fully commit the retirement decision (but not the consumption decision) in period 1 and self 2 cannot change it. Then self 1 would simply choose to work in period 2 if s_1^{*w} lead a higher life-time utility than s_1^{*r}. In reality, self 2 is not forced to work or to retire. This means that because of the concavity of $u(\cdot)$ self 1's optimal choice is one of the following: s_1^{*r}, s_1^{*w} or $\bar{k}_2$. This can better be seen from Figure 1, which shows the lifetime utility of self 1 as a function of s_1 for the two cases when self 2 works (U^w) and when he retires (U^r). In practice, self 1 would only consider the portion of U^w corresponding to $s_1 \leq \bar{k}_2$ and the portion of U^r corresponding to $s_1 > \bar{k}_2$ (he would not consider the dotted portion of the utility functions in the figure).

A problem could arise if the two optima lie on the same side with respect to $\bar{k}_2$, as it is the case in the figure. If $\bar{k}_2 < s_1^{*w} < s_1^{*r}$, then self 1 would have to choose between $\bar{k}_2$ and s_1^{*r} even though s_1^{*w} gives a higher utility. This is the case in the figure, where self 1's first best choice is to save s_1^{*w} and to work, but self 2 would retire for this level of savings. If $\bar{k}_2$ leads to the highest expected discounted utility then it means that self 1 is under-saving in order to induce self 2 to work since he cannot reach his first best choice s_1^{*w}. Self 1 would then prefer to induce working in period 2, which was his first best choice regarding the retirement decision, but he has to give up his optimal saving plan. On the other hand, if s_1^{*r} gives a higher utility than $\bar{k}_2$, even if lower than s_1^{*w}, self 1 is over-saving to compensate for the early retirement decision of self 2.

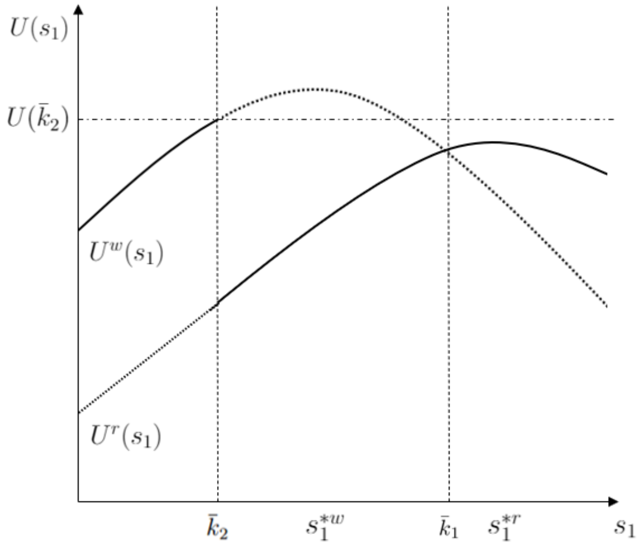

Figure 1: Life-time utility of self 1

Result 4. *Quasi-hyperbolic discounting might induce a sophisticated agent to over or under-save compared to his first best choice. In particular, a sophisticated agent might under-save to induce his future self to work, or might over-save to compensate for his future self's decision to retire early.*

5 Commitment Device

Consumption Decision

As in the previous section, I start discussing the consumption choice alone and then add the retirement decision in the model. Assume now that an illiquid asset is made available in the assets market. I assume in this section that the agent is sophisticated since a naïve agent would not realize that it might be useful to constraint his future behaviour with an illiquid asset. The asset is illiquid in the sense that if you buy it in period t it yields a return in period $t+2$, and not in period $t+1$ as the liquid asset. Also, if you decide to sell the illiquid asset in period t you would get paid in period $t+1$ because it is illiquid. If a consumer applies for a loan at time period t, the associated cash flow will not be available for consumption until time period $t+1$. These assumptions are the same as in Laibson (1997). In period 1, self 1 can decide to invest a share α of his savings in the liquid asset and $(1-\alpha)$ in the illiquid one, which yields the same return R as the liquid one.[1] In period 2, self 2 will have at his disposal only the part of savings that has been invested in the liquid asset: $\alpha s_1 R$. In period 3, self 3 will consume what is left, that is $(\alpha s_1 R - c_2)R$ plus $(1-\alpha)s_1 R^2$.

One of the most debatable features of Laibson (1997)'s model is that on the optimal path consumption is always larger than work salary. This feature is what leads illiquid asset to be a perfect commitment device in his multi-periods model. In every period t, the salary w_t is not enough for self t to consume his preferred quantity $c_t > w_t$. Therefore, he has to rely on the savings accumulated earlier s_{t-1}. Previous-selves would like self t to consume less than he would like, say $\hat{c}_t$ (with $\hat{c}_t < c_t$). Therefore, they can simply invest a share $(1-\alpha)$ of the savings in the illiquid asset such that $\alpha s_{t-1} + w_t = \hat{c}_t < c_t$. In this way previous selves can fully commit self t's consumption choice.

[1] This is in line with Laibson (1997). The qualitative results do not depend on the identical returns assumption.

There are different reasons why consumption could be larger than salary. First, we could assume that initial wealth w_0 is high enough but not model how this wealth is accumulated (e.g. in previous periods). Second, we could think that self 1 massively borrowed in the initial period. Third, we can imagine a situation where period 1 is the last period when the agent is working and will be forced to retire in period 2. The discussion below refers to this last scenario.

The maximization problem of self 2 in now slightly different, because he cannot consume his preferred quantity c_2 if this is larger than the liquid wealth he can spend. This is formalized as:

$$\max_{c_2,c_3} u(c_2) + \beta\delta u(c_3) \tag{39}$$

$$\text{s.t. } c_2 \leq \alpha s_1 R \tag{40}$$

$$c_3 = (\alpha s_1 R - c_2)R + (1-\alpha)s_1 R^2 \tag{41}$$

$$\implies c_2^{**} = \begin{cases} \frac{1}{1+\beta\delta} s_1 R & \text{if } \frac{1}{1+\beta\delta} s_1 R \leq \alpha s_1 R \\ \alpha s_1 R & \text{if } \frac{1}{1+\beta\delta} s_1 R > \alpha s_1 R \end{cases} \tag{42}$$

$$c_3^{**} = \begin{cases} \frac{\beta\delta}{1+\beta\delta} s_1 R^2 & \text{if } \frac{1}{1+\beta\delta} s_1 R \leq \alpha s_1 R \\ (1-\alpha)s_1 R^2 & \text{if } \frac{1}{1+\beta\delta} s_1 R > \alpha s_1 R \end{cases} \tag{43}$$

The solution of the problem in (42) simply formalizes the intuition that self 2 can consume his preferred quantity if he has enough liquid wealth to buy it; if the liquid wealth at his disposal is not enough he would just consume all of it to get as close as possible to his preferred quantity (this directly follows from the concavity of $u(\cdot)$). This also shows how self 1 can constrain the choice of self 2 (c_2^{**}, c_3^{**}) through the parameter α. I have already shown in (20) how self 1 would like to consume $c_2^* = \frac{1}{1+\delta} s_1 R$ in period 2 and his optimal saving is given by (26). With logarithmic utilities, the choice of c_1/s_1 is not affected by the choice of c_2 and c_3. This means that if self 1 can influence the consumption choice of self 2 he still chooses the same c_1/s_1 compared to when he cannot influence self 2 choice. Therefore, what the illiquid asset does is to shift upward the life-time utility of the agent. Self 1 can set $\alpha = \frac{1}{1+\delta} < \frac{1}{1+\beta\delta}$ and $s_1 = s_1^*$. In this way, self 2 will be forced to consume all the liquid asset in period 2 and all the illiquid asset in period 3, where these consumption levels are exactly the first best solution of self 1. Therefore, in this setting, a sophisticated agent can perfectly commit his future selves by simply using an illiquid asset. Demand for the commitment device, however, is then a signal of sophistication only if consumption is larger than working salary. The following result simply mimics Laibson (1997)'s finding in a different model:

Result 5. *A sophisticated agent can fully commit his future consumption behaviour using an illiquid asset as a commitment device if on the optimal path consumption is always larger than working salary.*

If optimal consumption is lower than working salary one way to effectively affect future selves is by directly committing them. For example, by signing a contract that prescribes that in each period a share of the future salary has to be invested in illiquid assets. This is something in the spirit of the SMarT program proposed by Thaler and Benartzi (2004), where people can commit in advance to allocating a portion of their future salary increases toward retirement savings. In order for the commitment to be effective, the available wealth should be enough to buy prior-selves optimal consumption quantity but not enough to buy current-self optimal quantity, or that the cost implied by changing consumption planning is higher than the extra utility.

From a theoretical perspective, the same goal of the SMarT program can be achieved by allowing $\alpha < 0$. Consider an agent who wants to consume $c_2 < w_2$, but his previous-self would like him to consume $\hat{c}_2 < c_2 < w_2$. This previous-self saved s_1 and can set α such that $\hat{c}_2 = w_2 + \alpha s_1$, which implies a negative value for α (if s_1 is positive). What the previous-self is doing is borrowing money (negative investment in liquid assets) against his future salary and using all the savings s_1 plus the borrowed money to invest in illiquid assets. At time 2, the agent is forced to repay his debt and he has to use part of his wage to pay it since all the savings are illiquid. Moreover, he cannot borrow against the illiquid assets to consume more in period 2 since such borrowing takes one period to be effective. Self 1 is basically committing in advance to allocating a part of his future salary towards savings for the last period.

Result 6. *A sophisticated agent can fully commit his future consumption behaviour using an illiquid asset.*

Retirement Decision

Let's now introduce the retirement decision in period 2. Self 2's allocation decision if he retires is the same as in the case where he is forced to retire and thus the solution is still the one in (42). We refer to this allocation as (c_2^{r**}, c_3^{r**}). Self 2's allocation problem if he decides to work is given by the following, with the liquidity constraint on consumption in period 2:

$$\max_{c_2,c_3} u(c_2) - e + \beta\delta u(c_3) \tag{44}$$

$$\text{s.t. } c_2 \le \alpha s_1 R + w_2 \tag{45}$$

$$c_3 = (\alpha s_1 R + w_2 - c_2)R + (1-\alpha)s_1 R^2 \tag{46}$$

$$\implies c_2^{w**} = \begin{cases} \frac{1}{1+\beta\delta}(s_1 R + w_2) & \text{if } \frac{1}{1+\beta\delta}(s_1 R + w_2) \le \alpha s_1 R + w_2 \\ \alpha s_1 R + w_2 & \text{if } \frac{1}{1+\beta\delta}(s_1 R + w_2) > \alpha s_1 R + w_2 \end{cases} \tag{47}$$

$$c_3^{w**} = \begin{cases} \frac{\beta\delta}{1+\beta\delta}(s_1 R + w_2)R & \text{if } \frac{1}{1+\beta\delta}(s_1 R + w_2) \le \alpha s_1 R + w_2 \\ (1-\alpha)s_1 R^2 & \text{if } \frac{1}{1+\beta\delta}(s_1 R + w_2) > \alpha s_1 R + w_2 \end{cases} \tag{48}$$

For given values of s_1 and α set by self 1, that self 2 takes as exogenous, self 2 will compute $(c_2^{r**}(s_1,\alpha), c_3^{r**}(s_1,\alpha))$ and $(c_2^{w**}(s_1,\alpha), c_3^{w**}(s_1,\alpha))$. He will then decide whether to work or not by comparing the utility deriving from the two different consumption plans. In this framework, the choice of self 1 regarding his savings (s_1, α) will influence both the consumption in periods 2 and 3 and the retirement decision. Self 1 problem is now more complicated. Simply setting α and s_1 to reach self 1's first best choice would not necessarily work, because this could change self 2 retirement decision in an unintended way. To gain intuition let's analyse the three possible different cases. In some cases self 2 would not work while self 1 would like him to, but the opposite cannot happen (from conditions (29) and (32)). Furthermore, self 1 would like to influence the allocation of consumption between periods 2 and 3. In particular, changing the allocation of consumption between periods 2 and 3 shifts the life-time utility of self 1 upwards or downward, without changing the saving level s_1 that maximizes the life-time utility. This means that self 1 can start by considering to stick to his first best savings level and use the illiquid asset to reallocate consumption from period 2 to period 3.

Case 1: Self 1 and self 2 agree that it is optimal to work in period 2 given self 1 optimal saving s_1^{*w} and no commitment. There is then no need to influence self 2 retirement decision. Self 1 could stick to s_1^{*w} and set α to influence self 2 consumption decision. In particular, self 1 sets α such that:

$$\frac{1}{1+\delta}(s_1^{*w} R + w_2) = \alpha s_1^{*w} R + w_2 \implies \alpha^* = \frac{\frac{1}{1+\delta}(s_1^{*w} R + w_2) - w_2}{s_1^{*w} R} \le 1 \tag{49}$$

The analytical solution is not bounded between zero and one:

- If $\alpha^* \in [0,1]$, then by construction of α^* self 2's constraint when working is binding because he has at his disposal $\frac{1}{1+\delta}(s_1^{*w} R + w_2) < \frac{1}{1+\beta\delta}(s_1^{*w} R + w_2)$. Note that self 2 would not be induced to change his behaviour and retire by the choice of α^*. This is because when the constraint in (47) binds so does the one in (42) and self 2 decision is still given by (29). Intuitively, if self 2 has less resources at his disposal (because of the liquidity constraint) he is even more motivated to work and not to retire. In this case, the agent can again fully commit himself.

- If $\alpha^* < 0$, this means that the wage in period 2 is so large that even when all the savings are illiquid ($\alpha = 0$) self 2 would consume more than self 1 would like him to. If we stick to Laibson (1997) conclusion that consumption is larger than labor income in every period then this case is ruled out. However, this is not the case for high salary in period 2 (or low initial endowment):

$$w_2 \geq \frac{s_1^{*w} R}{\delta} \text{ with } s_1^{*w} = \frac{\beta\delta^2(w_0 + w_1) + \beta\delta(w_0 + w_1) - w_2/R}{\beta\delta^2 + \beta\delta + 1} \tag{50}$$

$$\implies w_2 \geq \frac{\beta\delta^2 + \beta\delta}{\delta(\beta\delta^2 + \beta\delta + 1) + 1} R(w_0 + w_1) \tag{51}$$

If this holds, $\alpha^* < 0$ implies that the agent is borrowing money to leave to his future-self a debt that has to be paid from his future working salary, thus decreasing his disposable income and constraining his consumption level. The agent is simply reallocating resources from period 2 to period 3 to smooth consumption. As above, self 2 would not be induced to change his behaviour and retire by the choice of α^* and also in this case the agent can fully commit himself.

Case 2: Self 1 and self 2 agree that it is optimal to retire in period 2 given self 1 optimal saving s_1^{*r} and no commitment. As before, we can start with the guess that self 1 would stick to s_1^{*r} and he would use α to reallocate consumption between periods 2 and 3, that is:

$$\frac{1}{1+\delta} s_1^{*r} R = \alpha s_1^{*r} R \implies \alpha^* = \frac{1}{1+\delta} \in [0, 1] \tag{52}$$

Given the lower liquid wealth at his disposal, however, self 2 might change his decision and be induced to work. Again, by construction self 2 liquidity constraint when retiring (42) is binding because $\frac{1}{1+\delta} s_1^{*r} R < \frac{1}{1+\beta\delta} s_1^{*r} R$. However, self 2 constraint when working (47) is not necessarily binding, which means that this level of α would not only induce Self 2 to work but could also be ineffective in affecting consumption.

If $\alpha^* = \frac{1}{1+\delta}$ does not induce working self 1 can reach his first best allocation and retirement decision. The illiquid asset is a perfect commitment device.

If $\alpha^* = \frac{1}{1+\delta}$ induces working self 1 optimal choice is given by one of the following:

- Set $s_1 = s_1^{*r}$, the level of savings that induces retirement, and set the level of α to the minimum level larger than α^* that does not induce working. There always exists such a level that is large enough to induce retirement but small enough to slightly improve the consumption allocation from the perspective of self 1. Self 1 induces his optimal decision about retirement but cannot induce his optimal consumption allocation. In this case the fact that the commitment pushes one action in an unintended direction implies a lower demand for commitment compared to the case where the agent is forced to retire, that is the agent would demand more of the commitment if he was addressing only one decision instead of two decisions.
- Set $s_1 = s_1^{*w}$, the level of savings that induces working, and set the level of α to reach his first best allocation for periods 2 and 3 ($\alpha = \frac{\frac{1}{1+\delta}(s_1^{*w} R + w_2) - w_2}{s_1^{*w} R} \lesseqgtr \alpha^*$, as in Case 1, which in fact would induce working together with $s_1 = s_1^{*w}$). Self 1 induces his optimal consumption allocation but cannot induce his optimal decision about

retirement, even if self 1 and self 2 agreed to retire given self 1's first best choice s_1^{*r}. In this case the fact that the commitment pushes one action in an unintended direction can imply both a higher or a lower demand for commitment compared to the case where the agent is forced to retire.

Case 3: Self 1 and 2 disagree on the retirement decision given self 1 optimal saving s_1^{*w} and no commitment. When there is time inconsistency regarding both consumption and retirement self 1 is in a better position to constrain his future self. In fact, in order to force self 2 to work self 1 has to set a low α. Similarly, in order to affect the consumption he has to set a low α. Self 1 could set $s_1 = s_1^{*w}$, which would actually induce retirement if no commitment is provided. His optimal consumption allocation is induced by the same α^* as in case 1 ($\alpha^* = \frac{\frac{1}{1+\delta}(s_1^{*w}R+w_2)-w_2}{s_1^{*w}R}$). If this value is low enough to induce working then full commitment can be achieved. If α^* is not low enough to induce working then self 1's optimal choice is given by one of the following:

- Set $s_1 = s_1^{*r}$ and set the level of α to induce the optimal consumption allocation between period 2 and 3 (with $\alpha = \frac{1}{1+\delta} \lesseqgtr \alpha^*$). Self 1 induces his optimal consumption allocation but cannot induce his optimal decision about retirement.
- Set $s_1 = s_1^{*w}$, which would induce retirement without the commitment, but use a large quantity of the illiquid asset ($\alpha < \alpha^*$) to force self 2 to work. Self 1 induces his optimal decision about retirement but cannot induce his optimal consumption allocation.

The possibility of buying an illiquid asset increases self 1's life-time utility if he is sophisticated and has to decide how much to consume and when to retire. However, if the agent is not supplying work inelastically, the illiquid asset may be not sufficient to reach a full commitment regarding the two decisions. In particular, a model where an agent only consumes can possibly over/under-estimate the demand for commitment if the agent is actually deciding also when to retire.

Result 7. *A sophisticated agent is not necessarily able to perfectly commit his future consumption and retirement behaviours using an illiquid asset.*

Partially Naïve Agents

I have assumed so far, in the discussion concerning self-commitment devices, that our representative agent is fully rational and sophisticated. Both these assumptions could rightly be questioned. In particular, it is uncertain whether time inconsistent people are aware of their bias. It is likely that there is great heterogeneity across individuals and perhaps also across different situations. It seems intuitive that people are more prone to recognize their self-control problems when they concern actions that are repeated frequently over time. Some classic examples would be smoking, alcohol consumption or gym attendance (see for example Giné et al. (2010), Schilbach (2019) and DellaVigna and Malmendier (2006)).

In our settings, it is not too restrictive to assume that individuals are aware of their present-biased preferences over consumption choices. Since this is an action that we repeat daily we tend to be aware of our time preferences over consumption. However, since the retirement decision is usually taken only once, there is less scope for learning. I then refer to partially naïve agents as to individuals who are aware of their present-biased preferences over consumption, but who ignore their bias over the retirement decision.

Fully naïve agents would clearly not exhibit a demand for commitment devices. The result is that, regardless of their retirement decision in period 2, they would end up with a sub-optimal consumption allocation between periods 2 and 3 from the perspective of self 1.

A partially naïve agent would instead always demand some amount of the illiquid asset to reallocate consumption from period 2 to period 3, regardless of his retirement plan. However, this choice could have an unintended effect on the retirement decision in period 2: self 1 could force self 2 to work even if self 1 wanted him to retire. Ultimately, this would make the use of the commitment device undesirable. In fact, if self 2 is induced to work he would also be able to break the liquidity constraint. Therefore, the illiquid asset would induce working, which was not optimal, without improving the consumption allocation. In that case, being partially naïve would result in a lower welfare compared to being naïve.

Mandatory Retirement Plan

In the previous section, I showed how illiquid assets can affect sophisticated people's choices. I used illiquid assets as an example of self-commitment device because they are popular in the literature, but also because the underlying mechanism is very similar to that of mandatory retirement plans. Most employed workers, in fact, cannot fully dispose of their working salary: part of it is illiquid and will be available only after retirement. The primary goal is to avoid that present-biased people under-save when they are young and regret it when they are old. This is the aim of the so-called Pillar 1 and 2, that is the pillars of pension schemes that address the risks of individual myopia and low earnings, which are usually a mandatory component of pension systems that depend on public contributions.

Moreover, as Laibson (1997) showed, illiquid assets are intuitively effective when optimal consumption is larger than working salary, which is usually the case, for example, when people are already retired. This is because when the time comes, future selves are impatient and post-pone savings. If they are exogenously forced to save, however, life-time utility can be increased. This means that not only naïve but also sophisticated agents could benefit from exogenously set commitments. Still, the commitment is effective only if the mandatory savings do not leave enough liquid wealth for the current-self to consume his preferred quantity. Moreover, illiquid assets commitment and mandatory retirement plans share another drawback: they can both induce inefficient labour supply. This is even more likely when the saving rate is exogenously imposed as a flat rate across individuals that adhere to the same retirement plan. Furthermore, as the case of partially naïve agents shows, mandatory savings plan could induce an undesirably high labour supply without improving the consumption allocation of the agent.

6 Conclusions

First, this paper presents a model that is able to generate time inconsistency regarding endogenous consumption and retirement decisions. The model is particularly simple and tractable, yet it provides intuition for the underlying mechanisms. It builds on existing works that, however, have been proved to be not fully satisfying in their conclusions as they did not formally proved some claim (Diamond and Koszegi, 2003) or because they proved it in particularly narrowed version of the model (Holmes, 2010). The result is that naïve agents with quasi-hyperbolic discounting functions could exhibit time-inconsistent behaviour regarding both retirement and consumption planning. Also, a sophisticated agent could be induced to over or under-save to affect his future behaviour.

Second, it provides evidence that illiquid assets can be used as a commitment device also to affect retirement decisions. This generalizes the intuition derived from Laibson (1997) and also extends the model of Diamond and Koszegi (2003) by introducing a market for illiquid assets.

Third, it shows that when two interdependent choices have to be taken, such as consumption and retirement, then demand for commitment is non-trivial. In this case, the effectiveness of one commitment device can be high or low depending on whether it can address both actions in the desired way at the same time. This could explain why the observed demand for commitment devices is usually lower than theoretical models with a single decision would predict, but also why some studies found large demand and even willingness to pay. In particular, if we consider a more realistic environment with uncertainty, risk aversion and many actions and consequences to be considered, I would expect even more mixed results depending on preferences, sophistication degree, financial literacy and other unobservable characteristics. In particular, it is easy to imagine that uncertainty would provide an incentive not to reduce the agent's own future set of actions.

Fourth, the demand for commitment is not strictly monotone in the sophistication degree. Partially naïve agents could demand more illiquid assets than sophisticated agents, because they do not realize that such a large quantity could push their retirement decision in an unintended direction. This means that by looking at commitment demand we cannot necessarily infer people's sophistication level. This also points against the use of a revealed preferences argument to infer the primitives of people utility from observational data.

All these considerations suggest caution in using observational data when the researchers have little control over the set of actions that an agent is considering at the same time. This would particularly apply in the context of lifetime consumption models and their calibration. At the same time, this paper suggests that such models should incorporate the retirement decision both to better understand the demand for illiquid assets and to inform the design of optimal mandatory saving plans that takes into account the effect of illiquid savings on the retirement choice.

References

Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines. The Quarterly Journal of Economics 121, 635–672. URL: https://doi.org/10.1162/qjec.2006.121.2.635, doi:10.1162/qjec.2006.121.2.635, arXiv:http://oup.prod.sis.lan/qje/article-pdf/121/2/635/5324429/121-2-635.pdf.

Augenblick, N., Niederle, M., Sprenger, C., 2015. Working over Time: Dynamic Inconsistency in Real Effort Tasks. The Quarterly Journal of Economics 130, 1067–1115. URL: https://doi.org/10.1093/qje/qjv020, doi:10.1093/qje/qjv020, arXiv:http://oup.prod.sis.lan/qje/article-pdf/130/3/1067/30637254/qjv020.pdf.

Beshears, J., Choi, J.J., Harris, C., Laibson, D., Madrian, B.C., Sakong, J., 2015. Self Control and Commitment: Can Decreasing the Liquidity of a Savings Account Increase Deposits? NBER Working Papers 21474. National Bureau of Economic Research, Inc. URL: https://ideas.repec.org/p/nbr/nberwo/21474.html.

DellaVigna, S., Malmendier, U., 2006. Paying not to go to the gym. American Economic Review 96, 694–719. URL: http://www.aeaweb.org/articles?id=10.1257/aer.96.3.694, doi:10.1257/aer.96.3.694.

Diamond, P., Koszegi, B., 2003. Quasi-hyperbolic discounting and retirement. Journal of Public Economics 87, 1839–1872. URL: https://EconPapers.repec.org/RePEc:eee:pubeco:v:87:y:2003:i:9-10:p:1839-1872.

Feigenbaum, J.A., Findley, T.S., 2013. Quasi-hyperbolic discounting and the existence of time-inconsistent retirement. Theoretical Economics Letters 3, 119–123.

Feigenbaum, J.A., Findley, T.S., 2015. Quasi-hyperbolic discounting and delayed retirement. Theoretical Economics Letters 5, 325–331.

Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: A critical review. Journal of Economic Literature 40, 351–401. URL: http://www.aeaweb.org/articles?id=10.1257/002205102320161311, doi:10.1257/002205102320161311.

Giné, X., Karlan, D., Zinman, J., 2010. Put your money where your butt is: A commitment contract for smoking cessation. American Economic Journal: Applied Economics 2, 213–35. URL: https://www.aeaweb.org/articles?id=10.1257/app.2.4.213, doi:10.1257/app.2.4.213.

Holmes, C., 2010. Quasi-hyperbolic preferences and retirement: A comment. Journal of Public Economics 94, 129–130.

Kaur, S., Kremer, M., Mullainathan, S., 2010. Self-control and the development of work arrangements. American Economic Review Papers and Proceedings 100, 624–28. URL: http://www.aeaweb.org/articles?id=10.1257/aer.100.2.624, doi:10.1257/aer.100.2.624.

Laibson, D., 1997. Golden Eggs and Hyperbolic Discounting. The Quarterly Journal of Economics 112, 443–478. URL: https://doi.org/10.1162/003355397555253, doi:10.1162/003355397555253, arXiv:http://oup.prod.sis.lan/qje/article-pdf/112/2/443/5291736/112-2-443.pdf.

Laibson, D., 2015. Why don't present-biased agents make commitments? American Economic Review Papers and Proceedings 105, 267–272. URL: https://www.aeaweb.org/articles.php?doi=10.1257/aer.p20151084.

O'Donoghue, T., Rabin, M., 1999. Doing it now or later. American Economic Review 89, 103–124. URL: http://www.aeaweb.org/articles?id=10.1257/aer.89.1.103, doi:10.1257/aer.89.1.103.

Phelps, E.S., Pollak, R.A., 1968. On second-best national saving and game-equilibrium growth. The Review of Economic Studies 35, 185–199. URL: http://www.jstor.org/stable/2296547.

Samuelson, P.A., 1937. A note on measurement of utility. The Review of Economic Studies 4, 155–161. URL: http://www.jstor.org/stable/2967612.

Schilbach, F., 2019. Alcohol and Self-Control: A Field Experiment in India. American Economic Review 109, 1290–1322. URL: https://ideas.repec.org/a/aea/aecrev/v109y2019i4p1290-1322.html.

Thaler, R., 1981. Some empirical evidence on dynamic inconsistency. Economics Letters 8, 201–207. URL: https://EconPapers.repec.org/RePEc:eee:ecolet:v:8:y:1981:i:3:p:201-207.

Thaler, R.H., Benartzi, S., 2004. Save more tomorrow™: Using behavioral economics to increase employee saving. Journal of Political Economy 112, S164–S187. URL: https://doi.org/10.1086/380085, doi:10.1086/380085, arXiv:https://doi.org/10.1086/380085.