Time-Inconsistency, Commitment, and Learning

Inaugural-Dissertation
zur Erlangung des Grades
Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2018

vorgelegt von

Else Gry Bro Christensen

Referent: Prof. Dr. Klaus M. Schmidt
Korreferent: Prof. Dr. Florian Englmaier
Promotionsabschlussberatung: 11. Juli 2018

Datum der mündlichen Prüfung: 26.06.2018

Namen der Berichterstatter: Klaus M. Schmidt, Florian Englmaier, Fabian Herweg

Acknowledgments

I would like to thank my supervisor, Klaus Schmidt, for giving me this opportunity and for challenging me to become a better researcher, and my second advisor, Florian Englmaier, for much advice and support, for including me, and always making me feel welcome at his chair. I would also like to thank my coauthors: Julia Nafziger and Takeshi Murooka without whom this dissertation would never have been possible. Furthermore, I would like to thank my colleagues, in particular Marco Schwarz, Dominik Fischer, and Maximilian Breu, for many talks and many lunches. I would also like to thank the organizers at the MGSE for always being helpful, the GRK 1928 with coordinator Carsten Eckel, and the CRC TR 190 "Rationality and Competition" for funding. Last but not least, I would like to thank my family and my friends for all their support and patience, you have been invaluable.

Contents

Preface

Self-control problems are a common challenge for many people. People give in to temptation in the moment and overindulge in unhealthy behavior or postpone unpleasant tasks against their own long-run interests. For example, they are tempted to overeat or overspend and to procrastinate unpleasant tasks such as dieting or working.

In this dissertation I model three different settings where present-biased preferences have important economic implications related to the following three questions: Why do people keep making the same plans and fail? Why do people not seek help to follow through on their plans? And how might firms respond to consumers with such behaviors?

The first chapter explores the consequences of present-bias in marked interaction between present-biased consumers and rational firms. The second chapter examines the scope for overcoming present-bias when agents can construct their own commitment devices. The third chapter proposes a mechanism for why agents may not even learn about their present-bias even when given ample learning opportunities. In the following, I give a summary of the papers and their findings.

Chapter 1: In this chapter, we analyze how firms can choose packaging sizes of tempting foods to exploit consumers present-bias or help them overcome their present-bias.

We consider the shopping and consumption decision of an individual with a self-control problem. The consumer believes that restricting the consumption of a sinful product (such as chips) is in his long-run interest. But when facing the actual decision he is tempted to overeat. The individual makes a shopping trip before he consumes, but he can also go shopping when he is tempted to overeat. We ask how firms react to such self-control problems, and possibly exploit them, by offering different package sizes. In a competitive market, either one or three package sizes are offered. In contrast to common intuition, the large, and not the small package might be a commitment device. The latter may serve to exploit the naive consumer.

A monopolist offers either one package size to cater to the interests of the long run self; or he offers a small, “commitment"-package for the sophisticated and a large package for the naive consumer. While the sophisticated consumer receives the same commitment in the monopolistic and competitive market, the social surplus from the naive consumer can be lower in the competitive market because the naive consumer goes shopping more often in competitive markets.

Chapter 2: In this chapter, I suggest a new explanation for why commitment uptake is rare.

I examine a setting where present-biased agents face stochastic costs and fixed benefits from completing a task. Agents can take up commitment to attempt to overcome their present-bias. The commitment devices consist of a fee that agents have to pay if they do not complete the task. The fee cannot be contingent on the cost realization but the devices are themselves free and the agent can choose any fee level. Nevertheless, I find that commitment is rarely optimal even for perfectly sophisticated agents. If the task is worthwhile for some cost realizations but the expected cost is higher than the benefit then agents prefer never completing the task over commitment that ensures the task will always be completed. Therefore, if the agent demands any commitment it is partial and the agent faces a tradeoff between completing the task only when worthwhile and incurring the fee with positive probability. This implies, that commitment can be socially wasteful.

Partial commitment is more likely to be optimal for high variance of costs, it is ambiguously related to the agent's (perceived) degree of present-bias. However, demand for partial commitment is decreasing in present-bias for severe present-bias, while overconfidence can increase the likelihood of taking up commitment.

Chapter 3: In this chapter, we propose an explanation for why present-biased agents do not learn about their present-bias over time.

PREFACE

We investigate models in which an agent aims to complete a task. While the agent is initially uncertain about own future self-control, in each period she has an opportunity to learn about own future self-control by incurring a non-negative (physical or psychological) cost. If the agent is time-consistent, she always chooses to learn whenever the learning is beneficial. If the agent has time-inconsistent preferences, however, she may procrastinate such a learning opportunity. Even when the cost of learning is zero, the agent may procrastinate if the learning takes time and her preference exhibits an intertemporal conflict between future selves (e.g., hyperbolic discounting). We also derive conditions in which procrastination of learning and non-completion of a task is a unique equilibrium outcome. When the agent has multiple initially-uncertain attributes (e.g., own future self-control and own ability for the task), the agent's endogenous learning decisions may be misdirected—she chooses to learn what she should not learn from the long-run perspective, and she chooses not to learn what she should.

Connection: The papers are tightly connected. They each speak to both the consequences of present-bias and the difficulties in addressing self-control issues.

Chapter 1 shows both the risks of being unaware of own present-bias (by showing how naive consumers can be exploited) and the caution policymakers must take if attempting to regulate markets (it is not necessarily the largest packages of unhealthy foods that induce overconsumption). The chapter considers a setting where commitment devices are not available and, therefore, is complemented by Chapter 2 which examines the scope for commitment. Though the setting described in Chapter 2 differs from that of Chapter 1 both models consider the same intra-personal conflict: the agent's preferences over a task or good are time-inconsistent. Chapter 2 considers investment goods rather than leisure goods and introduces stochastic costs. In this setting, even sophisticated agents may not benefit from commitment, and therefore, may still incur negative welfare consequences of present-bias. Chapter 3 bridges the two other chapters by addressing how agents can stay naive. It thereby shows how sophisticated and naive agents can coexist and therefore why the behavior of both types are relevant to analyze.

Chapter 1

Packaging of Sin Goods - Commitment or Exploitation?*

1.1 Introduction

Many individuals face self-control problems. The pleasure of the moment seduces them to act against their own long run interests. For example, they are tempted to shirk on unpleasant tasks – such as dieting. And a poor diet contributes to the problem of overweight and obesity. The World Health Organization reports that more than 1.4 billion adults were overweight in 2008, and more than half a billion obese. It estimates that at least 2.8 million people die each year as a result of being overweight or obese. Moreover, globally, 44 percent of diabetes, 23 percent of ischaemic heart disease and 7-41 percent of certain cancers are attributable to overweight and obesity.[1] The associated health costs are large (cf., e.g., Finkelstein et al., 2009).

Chandon and Wansink (2010) discuss how firms influence food intake with their marketing strategies and thereby may contribute to the problem of overeating. Examples of such marketing strategies are food prices and promotion, the food's quality and quantity, marketing, the availability, salience and convenience of food, the type, size and shape of serving containers, or the atmospherics of the purchase and consumption environment.

---

* This chapter is based on joint work with Julia Nafziger which has already been published under the same title in the Journal of Economic Behavior & Organization.

[1] See *http://www.who.int/features/factfiles/obesity/en/index.html* (last accessed January 2014).

In this paper, we want to focus on one particular marketing strategy – the packaging of sinful products such as chips. Wertenbroch (1998) argues that consumers engage in pre-commitment by rationing their purchase quantities, i.e., by buying, for example, small packages. Two questions arise from this. First, can the consumer indeed limit his consumption through such a strategy? And second, how do firms react to the consumer's self-control problem? The argument by Wertenbroch (1998) presumes that firms indeed offer small packages as commitment devices. But do they indeed do so or are they trying to counteract the consumer's wish for commitment?

To answer these questions, we consider the shopping and consumption decision of a vice, or sinful good (such as chips, cigarettes, or chocolate) of an individual that faces a self-control problem that arises due to time-inconsistent preferences. The individual judges that limiting the amount of, say, chips consumption is in his long-run interest. But once he sits in front of the TV and starts eating chips, the distant health benefits of a healthier life style suddenly do not seem worth the effort of restricting consumption.

The consumer goes shopping when he is not tempted to overeat. For example, when doing his weekly shopping trip, the consumer (self 0) has planned beforehand how much chips to buy and is not hungry. When sitting in front of the TV in the evening, however, the consumer (self 1) is tempted to overeat chips. The consumer can go shopping at this point, but, because of opportunity costs, the costs of such spontaneous shopping trips are higher than those of his weekly, planned shopping trip. We assume that the consumer is either sophisticated or naive, which means that the consumer, when doing his weekly shopping trip, is either fully aware or not at all aware that he faces a self-control problem.

Firms offer the consumer to buy a certain quantity (a "package") for a transfer. In the main model, we consider a competitive market. The sophisticated self 0 perfectly anticipates the shopping and consumption decision of his future self. Firms respond by offering self 0 either full or partial commitment. More precisely, if the the shopping costs of self 1 are large, self 0 buys the package that is optimal from his viewpoint. In this case, self 1 is not tempted to go shopping again given the package self 0 bought and so the market can provide full-commitment. If, however, self 1 would be tempted to go shopping again if self 0 bought the from his viewpoint optimal package, then only partial commitment is feasible. That is, firms offer a larger package to self 0 than is optimal from his perspective. Yet, self 0 buys this larger package because he knows that otherwise self 1 would go shopping again to buy a top-up. In contrast to the sophisticated self 0, the naive self 0 does not anticipate the decision of his future self. Hence, he goes shopping and buys the package that is optimal from his point of view. If the shopping costs of self 1 are sufficiently large, then self 1 will not go shopping. So in this case, self 1 will consume the package that is optimal from the point of view of self 0, i.e., the naive consumer receives full commitment. If, however, the shopping costs of self 1 are not sufficiently large, self 1 will go shopping again and buys a “top-up" package. So in this case, commitment breaks down.

The result yields several implications and testable predictions. First, we predict that commitment is easier to achieve if self 1 faces large shopping costs. There is some evidence that supports this result. Hinnosaar (2012) predicts that Sunday sales restriction should decrease weekend consumption of alcohol. When looking at an actual policy change, Bernheim et al. (2012) however observe no such effect. Currie et al. (2010) observe that a close geographical proximity of a fast food restaurant is associated with higher rates of obesity (of children and pregnant women). Leung et al. (2011) show that the availability of convenience stores within a close distance of residence is correlated with a greater risk of girls becoming overweight or obese. Lee (2012) however finds conflicting evidence regarding the association between distance and overweight.

Second, the result predicts that in the competitive market, if the shopping costs of self 1 are large relative to those of self 0, one package size is offered, which is tailored to the interests of self 0. If the shopping costs at time 1 are relatively small, three different package sizes (a small one, a medium one, and a large one) are offered, and thus the market does not offer commitment to the naive consumers. Consistent with this result Steenhuis et al. (2010) observe that firms offer different package sizes. In contrast to common intuition, we demonstrate that the small package is not necessarily a commitment device, but can serve to exploit the naive consumer, while the large package may offer commitment to the sophisticated consumers. Indeed, for empirically reasonable values of the self-control problem, the smallest package is always the exploitative.

Third, in a competitive market, the naive consumer may go inefficiently often shopping. He initially buys a relatively small package that is optimal from the point of view of self 0 – believing that he will only consume this package, i.e., not go shopping again. But later, when in the "hot state", he buys another top-up package. Indeed, Hinnosaar (2012) observes that time-inconsistent consumers go shopping more often than time-consistent consumer. And Vermeer et al. (2011) provide field evidence showing that people having a smaller portion in the lunch cafeteria later buy more other food.

We contrast the competitive market with a monopolistic one. When facing a naive consumer, the monopolist tailors the package to the preferences of the self from whom he can extract the highest surplus. Thus, if the shopping costs of self 1 are large relative to the shopping costs of self 0, he caters to self 0 and perfect commitment is possible. If not, he caters to the interests of self 1 and offers a relatively large package. The sophisticated self 0 is willing to pay for a smaller commitment-package. The monopolist offers such (partial or full) commitment products to the sophisticated consumer – possibly at a higher price. Thus, in a monopolistic market, small packages are always commitment devices. So overall, our results make clear that a careful market analysis and assessment of the self-control problem are needed to assess whether small packages are commitment devices or are exploitative.

Comparing the monopolistic to the competitive market shows that for the sophisticated consumer only the distribution of rents differs, but the social surplus is the same in both markets. In contrast, the naive individual might be better or worse off in the competitive market. On the one hand, the competitive market provides more often full commitment to the naive consumer than the monopolistic market. However, if he does not receive full commitment, the social surplus in the competitive market can be lower than in the monopolistic market because the naive consumer goes shopping too often in the competitive market. Thus, competition can decrease the social surplus.

Finally, we consider some robustness checks and extensions to our main analysis. We show that our results are robust to relaxing some of the more technical assumptions and robust to allowing for partial naïveté. We also extend the model – allowing for imperfect competition and market entry. In this setting, firms operate on two separate markets: one that caters self 0 and one that caters self 1. This can be seen as reflecting the difference between offering a product in standard supermarkets, or in late night convenience stores.

Related theoretical literature The theme that firms provide commitment to the sophisticated consumer, but exploit the naive consumer is well established in this literature on contracting with time-inconsistent agents. Our paper is distinct by the application and the repeated setting it studies: package sizes in consumer markets with re-shopping possibilities have been underexplored – and some new features arise in such a setting.

DellaVigna and Malmendier (2004) consider a model in which firms offer two-part tariffs consisting of a lump-sum payment and per-unit price to a time-inconsistent consumer. They establish that firms price investment goods below marginal costs, and leisure goods above. Gottlieb (2008) relaxes the assumption of DellaVigna and Malmendier (2004) that a consumer deals exclusively with one firm (for a similar model see also Kőszegi, 2005). The assumption of a competitive spot market is more realistic in, e.g., markets for consumption goods such as sinful products. He shows that in this case marginal cost pricing of sinful products arises, i.e., commitment for the sophisticated consumers vanishes. Our paper demonstrates that the market may provide commitment under some circumstances – to both sophisticated and naive consumers. Gottlieb (2008) also shows that competition can decrease the social surplus. In our model, this is the case for the naive, but not for the sophisticated consumer. Heidhues and Kőszegi (2010) consider a competitive credit market. They show that firms exploit naive consumers by offering a cheap baseline repayment, and specifying large penalties for late payments.[2]

A small literature strand in marketing asks about the package sizes firms provide when consumers face self-control problems. Dobson and Gerstner (2010) show why it can be profitable for firms to offer so-called “super-size" portions, i.e., very large portions which are not much more expensive than the normal sized portion. Firms employ such a strategy

[2] A number of papers focuses on the question how to screen agents who differ in their degree of sophistication, or in their degree of time-inconsistency (see, e.g., Eliaz and Spiegler, 2006, 2008; Esteban and Miyagawa, 2006; Esteban et al., 2007; Galperti, 2012). The theme that relatively sophisticated types receive full commitment, while relatively naive types are exploited re-appears when firms screen agents who differ in their degree of sophistication (see, e.g., Eliaz and Spiegler, 2006, 2008).

in order to price-discriminate between disciplined and tempted consumers. The former are willing to pay a premium for smaller sized portions. Jain (2012) considers, similar to our setting, the shopping decision of an individual with a present bias. The main difference is that in his framework only self 0 can go shopping, and consumption occurs on two days. The package size and number of package sizes are exogenous in his setting, while it is endogenous in ours. A small package thus is by assumption a commitment device in his model and its introduction increases the social surplus. Firms introduce small packages if the gain from doing so (attracting consumers who would otherwise not buy or buy less) outweighs the loss (some consumers buy less). Jain (2012) also briefly considers naive consumers, but as only self 0 can go shopping, exploitation by firms of naive consumers is not an issue.

Hinnosaar (2012) builds up a model of the behavior of a time-inconsistent consumer related to the one considered here. Her main aim is to identify time-consistent and time-inconsistent consumers from dynamic purchasing behavior, but she does not disentangle naive and sophisticated consumers. And she does not consider firm behavior as we do in this paper.

The paper is organized as follows. We first introduce the model in section 2.3. The main analysis and results are presented in section 1.3. In this section, we first consider the competitive market and then the monopolistic one. Section 1.4 discusses robustness and extensions. Section 1.5 concludes the paper.

1.2 Model

There is a continuum of firms who operate in a competitive market and a continuum of consumers (with unit mass one).[3] In each period $\tau \in \{0, 1\}$, firms offer the consumer a schedule, i.e., a quantity-transfer pair (x_τ, t_τ) that specifies for every quantity $x_\tau \in \mathbb{R}_0^+$, a (possibly negative) transfer $t_\tau \in \mathbb{R}$ from the consumer to the firm. Firms make these

[3] The assumption that there is a continuum of consumers is not crucial for the baseline model, but ensures that models in which the distribution of types of consumers play a role (see section 1.4) can easily be build on top of the existing framework.

	Type	True preferences	Beliefs about preferences at $\tau = 0$	Actions
$\tau = 0$	Naive Sophisticates	$u(x)$ $u(x)$	$u(x)$ $u(x)$	Shop $x_0 \geq 0$ if $x_0 > 0$: incur cost k_0
$\tau = 1$	Naive Sophisticates	$v(x)$ $v(x)$	$u(x)$ $v(x)$	Shop $x_1 \geq 0$ if $x_1 > 0$: incur cost k_1 consume $x_0 + x_1$

Table 1.1: *Overview of types, preferences, beliefs and actions*

offers in a given period simultaneously. In each period, firms observe which contracts have been offered previously. A firm's cost of producing x units is $c\,x$. The profits of the firm are

$$\Pi(x_0, x_1, t_0, t_1) = t_0 + t_1 - c\,(x_0 + x_1).$$

In period 0, the consumer has the opportunity to go shopping. In period 1, the consumer can go shopping again and consumption takes place. Shopping trips are costly. The consumer incurs a monetary cost k_τ for a shopping trip in period τ. We assume that $k_1 \geq k_0 = 0$. The timing and consumers' beliefs are summarized in table 1.1.

The consumer faces a self-control problem, i.e., the consumer's preferences regarding consumption change between periods 0 and 1. The date 0 incarnation of the consumer will be called self 0, and the date 1 incarnation will be called self 1. In period 0, the consumer's utility from consuming some quantity x is $u(x)$, while in period 1 the preferences have changed to $\nu(x)$. The consumer has quasi-linear preferences over consumption and the transfer/shopping costs. That is, self 0's utility function is given by:

$$U(x_0, x_1, t_0, t_1, k_1) = u(x_0 + x_1) - t_0 - t_1 - k_1.$$

And the one of self 1 is:

$$V(x_0, x_1, t_1, k_1) = v(x_0 + x_1) - t_1 - k_1.$$

Both $u(x)$ and $\nu(x)$ are strictly increasing and concave functions, implying $u'(x), v'(x) > 0$ and $u''(x), v''(x) < 0$. Further, $v(0) = u(0) = 0$. If the consumer does not consume, he receives the reservation utility $\bar{u} = 0$.[4] [5]

From the perspective of self 0, let the surplus maximizing quantity[6] be given by:

$$x_0^* = \arg\max_x \quad u(x) - c\,x,$$

and from the perspective of self 1:

$$x_1^* = \arg\max_x \quad v(x) - c\,x.$$

We are interested in the case where the good is a harmful vice good, such as chips, or cigarettes. Self 0 prefers a lower consumption than self 1. Assuming $v(x) > u(x)\,\forall x$ and $v'(x) > u'(x)\,\forall x$ implies $x_1^* > x_0^*$, i.e., there is a conflict of interest between self 0 and self 1.

Consumers can be either naive or sophisticated about their future preferences. The consumer believes that his future preferences are captured by the utility function $\hat{v}(x)$. If $\hat{v}(x) = v(x)$ the consumer is fully sophisticated and if $\hat{v}(x) = u(x)$ he is fully naive.[7] Partial naïveté does not change the main insights and is briefly discussed in section 1.4).

We assume that firms know $u(x)$, $v(x)$ and $\hat{v}(x)$. Relaxing the assumption that firms know the degree of sophistication does not change our results as we discuss in section 1.4. Furthermore, we assume that gains from trade are positive, i.e., $v(x_1^*) - c\,x_1^* - k_1 > 0$ and $u(x_0^*) - c\,x_0^* > 0$.

---

[4] Allowing for $\bar{u} > 0$, or $k_0 > 0$ does not change results. What matters is the difference between $\bar{u} + k_1$ and $\bar{u} + k_0$.

[5] These preferences capture in a stylized way present biased preferences (Laibson, 1997; O'Donoghue and Rabin, 1999). For example, the consumption of x units of the good causes immediate benefits of $b(x)$ and delayed costs of $\kappa(x)$. Self 0 weighs these costs equally, i.e., $u(x) = b(x) - \kappa(x)$. Self 1 attaches, due to his present bias, $\beta \in (0,1)$, a larger relative weight to the current benefits than to the delayed costs, i.e., $v(x) = b(x) - \beta\,\kappa(x)$.

[6] That is, the quantity that maximizes $U(x_0, x_1, t_0, t_1, k_1) + \Pi(x_0, x_1, t_0, t_1)$.

[7] We have a further implicit assumption; that consumers do not infer anything about own (changing) preferences from the quantity-transfer contracts on the market.

The setup captures the idea that in period 0 the individual does his (weekly) shopping trip. At this date, he is in a cold state and not tempted to overeat the harmful product (such as chips). In period 1, when sitting in front of the TV, he is in a hot state, and prefers to eat more chips than is optimal from the long run perspective. He can go shopping at this date, but at this point in time the opportunity costs of a shopping trip are higher than for the weekly shopping trip. For example, buying chips in the shop around the corner just before the movie starts causes higher costs than buying them along with other goods on a planned, weekly shopping trip.[8]

Competitive equilibrium As usual in the literature, we define the competitive equilibrium in terms of the contracts that survive competitive pressure. A contract is a quantity-transfer pair. Consumers apply for at most one contract and if their participation constraint is satisfied, they choose the contract that yields the highest utility to them (if indifferent they randomize 50-50). In each period, they can choose a different contract from a different firm. Each equilibrium contract earns zero expected profits, and there exists no profitable deviation in any period that is accepted by a consumer and that yields strictly positive expected profits. Each contract offered is purchased by some consumers. We assume that firms produce on the spot. Thus, firms can react to a deviation of a firm in the current period in later periods. In section 1.4, we discuss the robustness of the results to this assumption.

1.3 Analysis

We start by analyzing the case of naive consumers and then sophisticated consumers in a setting of perfect competition before contrasting these results to the case of a monopoly.

1.3.1 Naive consumers

Suppose self 0 bought some quantity x_0. Self 1 has to decide whether he is satisfied with this quantity, or whether he wants to incur the costs of an additional shopping trip and

[8] The assumption that costs are higher in period 1 than in period 0 could also reflect psychological costs, such as a bad conscience for an additional shopping trip.

pay the transfer to get the additional quantity x_1. So he goes shopping whenever the following participation constraint is satisfied:

$$\text{PC}_1 \quad v(x_0 + x_1) - k_1 - t_1 \geq v(x_0). \tag{1.1}$$

Since it will often be of interest when the participation constraint of self 1 is not satisfied, we will use the notation (NPC_1) when referring to the constraint that self 1 will not go shopping. The naive self 0 believes that his consumption plan coincides with the one of self 1. This means that self 0 does not anticipate that self 1 will possibly go shopping. Hence, the naive self 0 goes shopping whenever the following participation constraint is satisfied:

$$\text{PC}_0^N \quad u(x_0) - t_0 \geq 0. \tag{1.2}$$

Contract for self 1 A firm's problem in period 1 is to maximize its profits subject to the participation constraint of the consumer.

$$\max_{t_1, x_1} \quad t_1 - c\, x_1 \quad \text{ s.t. } (PC_1) \quad v(x_0 + x_1) - k_1 - t_1 \geq \max\{v(x_0), \bar{v}^*\}. \tag{1.3}$$

Competition affects the outside option of the consumer. If he does not buy from a specific firm, he can buy from another firm. Thus, in a competitive market, the outside option is determined endogenously. Specifically, utility level $\bar{v}^*$ is the perceived utility from the (equilibrium) contract. Since we are in a competitive setting, the maximized profit of each firm must be zero, i.e., $t_1 = c\, x_1$. In the optimization problem the level $\bar{v}^*$ ensures that this zero profit condition holds: $v^* = v(x_0 + x_1) - k_1 - t_1 = v(x_0 + x_1) - k_1 - c\, x_1$ for the equilibrium quantity x_1.

In the appendix, we show that the solution to this problem is to sell $x_1 = x_1^* - x_0$ at $t_1 = c\,(x_1^* - x_0)$, whenever $v(x_1^*) - c\, x_1^* - k_1 \geq v(x_0) - c\, x_0$, and 0 otherwise. When selling $x_1 = x_1^* - x_0$ to self 1, he will consume his preferred quantity $x_0 + x_1 = x_1^*$. That is, the utility of self 1 is maximized. The transfer ensures that firms make zero profits. The condition $v(x_1^*) - c\, x_1^* - k_1 \geq v(x_0) - c\, x_0$ is satisfied whenever self 1 finds it optimal to buy the equilibrium quantity at the respective transfer.

Contract for the naive self 0 We first argue that in any competitive equilibrium the naive self 0 goes shopping. If this were not the case, a firm had an incentive to deviate and offer self 0, say, x_0^* defined by $u'(x_0^*) = c$ for a transfer $t_0 = u(x_0^*)$. The naive self 0, who is not aware that self 1 shall possibly go shopping, accepts such an offer and the offer yields strictly positive profits for a firm. Thus, in a competitive equilibrium, the participation constraint of the naive self 0 cannot be violated and the period-0 contract for the naive consumer solves

$$\max_{t_0,x_0} \quad t_0 - c\,x_0 \quad \text{s.t.} \quad (PC_0^N) \quad u(x_0) - t_0 \geq \bar{u}^*. \tag{1.4}$$

Again, utility level $\bar{u}^*$ is the perceived utility from the equilibrium contract. The solution is to offer $(x_0^*, c\,x_0^*)$ in period 0. The naive self 0 accepts this offer.

1.3.2 Sophisticated consumers

In contrast to the naive, the sophisticated self 0 perfectly anticipates the consumption and shopping decision of his future self. We hence use a backward induction approach to derive the participation constraints. Suppose self 0 does not go shopping in period 0. Then self 1 would accept (x_1, t_1) if $v(x_1) - k_1 - t_1 \geq 0$. More generally, the participation constraint of sophisticated and naive consumers coincide at date 1, i.e., (PC_1) is also the relevant constraint for the sophisticated self 1. Contract (x_1, t_1) yields utility $u(x_1) - t_1 - k_1$ to self 0. Hence, the sophisticated self 0 accepts contracts which satisfy the following participation constraint

$$\text{PC}_0^S \quad u(x_0) - t_0 \geq u(x_1) - t_1 - k_1, \tag{1.5}$$

where (x_1, t_1) must be such that the participation constraint of self 1, (PC_1) (with $x_0 = 0$) holds and where x_0 must be such that self 1 does not go shopping. This means that for any profitable (x_1', t_1') it must hold that

$$\text{NPC}_1 \quad v(x_0) \geq v(x_0 + x_1') - k_1 - t_1'. \tag{1.6}$$

The intuition is as follows: The sophisticated consumer will know of the change in preferences, and therefore never buy a quantity x_0 so small, that there exists a quantity-transfer pair (x_1', t_1') such that self 1 will go shopping. Such a behavior would yield to additional shopping costs which the sophisticated self 0 could avoid by buying the total desired quantity immediately. Thus, (NPC_1) cannot be violated. But at the same time self 0 will never buy a quantity x_0 if he would have a higher utility from the expected equilibrium contract at date 1. This is what (PC_0^S) states.

Contract for self 1 While in equilibrium the sophisticated self 1 will never go shopping, we have to specify which contract firms offered him if he, off the equilibrium-path, were to go shopping. As the participation constraint for the naive and sophisticated consumer coincide in period 1, the optimization problem of a firm is identical to the one for the naive consumer, i.e., to (1.3). Again the solution to this problem is to sell $x_1 = x_1^* - x_0$ at $t_1 = c\,(x_1^* - x_0)$, whenever $v(x_1^*) - c\,x_1^* - k_1 \geq v(x_0) - c\,x_0$, and 0 otherwise.

Contract for the sophisticated self 0 Self 0 anticipates that if self 1 went shopping, he would consume x_1^* in total. Selling self 0 such a low quantity that self 1 goes shopping could not, however, be a competitive equilibrium. As the shopping costs of self 0 are lower than those of self 1, firms could increase profits by selling the desired quantity to self 0 rather than to self 1. Thus, in any competitive equilibrium, only the sophisticated self 0, not self 1 goes shopping. So the competitive equilibrium is characterized by the solution to the following optimization problem:

$$\begin{aligned} &\max_{t_0, x_0} \quad t_0 - c\,x_0 && (1.7)\\ &\text{s.t.} \quad (PC_0^S) \quad u(x_0) - t_0 \geq \bar{u}^*, \\ &\text{s.t.} \quad (NPC_1) \quad v(x_0) \geq v(x_1^*) - c\,(x_1^* - x_0) - k_1. \end{aligned}$$

Utility level $\bar{u}^*$ is the perceived utility from the equilibrium contract. Suppose first that $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$. Then self 1 would never go shopping when self 0 bought x_0^*, and hence the solution to the above problem is x_0^*, defined by $u'(x_0^*) = c$, $t_0 = c\,x_0^*$. Suppose next that self 1 would go shopping if self 0 bought only x_0^*. Then the "no-shopping" constraint, (NPC_1^S), has to be binding. Otherwise a firm had an incentive to deviate and lower x_0 by an ε. The lower offer is more attractive for self 0, and self 1 would, for ε small, still not go shopping. Hence, the optimal solution x_0^S is defined by $v(x_0^S) - c\,x_0^S = v(x_1^*) - c\,x_1^* - k_1$, and $t_0 = c\,x_0^S$. So overall, firms offer (partial) commitment to the sophisticated self 0. Specifically, selling self 0 the quantity desired by self 1, x_1^* for $t_1 = c\,x_1^*$, cannot be a competitive equilibrium. Suppose it were. Then a firm could deviate and offer self 0 some quantity $x_0 \in [x_0^*, x_1^*)$ and a transfer t_0, such that $u(x_0) - t_0 \geq u(x_1^*) - c\,x_1^*$. Such an offer raises the utility of self 0 if the quantity is such that self 1 would not go shopping in period 1. Consider, e.g., $x_0 = x_1^* - \varepsilon$ and note that $v(x_1^*) > v(x_1^*) - t_1 - k_1$. Offering $x_0 = x_1^* - \varepsilon$ implies that, by continuity, $v(x_1^* - \varepsilon) > v(x_1^*) - t_1 - k_1$.

Proposition 1.1.

1. *The competitive equilibrium contract for the naive individual is* $(x_0^*, c\,x_0^*)$*, which he buys in period 0. If, in addition,* PC_1 *is satisfied:* $v(x_1^*) - c\,x_1^* - k_1 \geq v(x_0^*) - c\,x_0^*$*, then firms offer contract* $(x_1^* - x_0^*, c\,(x_1^* - x_0^*))$*, which self 1 buys in period 1.*
2. *The competitive equilibrium contract for the sophisticated individual is* $(x_0^S, c\,x_0^S)$*, where* $x_0^S \in [x_0^*, x_1^*)$ *is defined by* $v(x_0^S) - c\,x_0^S = v(x_1^*) - c\,x_1^* - k_1$*, which he buys in period 0. If* $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$*, then* $x_0^S = x_0^*$*.*

In the competitive market, either 1 or 3 different package sizes are offered. If $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$, then both the naive and the sophisticated consumers receive full commitment and only a package of size x_0^* is offered. If this does not hold, then 3 packages are offered: a large package, a medium package and a small package. In the appendix, we provide some details on the ranking of the packages sizes. We demonstrate that for empirical reasonable values of the self-control problem the ranking is in contrast to common intuition. Specifically, for such values the largest package is x_1^S, i.e., the package that offers commitment to the sophisticated consumer. The small package is $x_0^* - x_1^*$ and thus not a commitment device, but serves to exploit the naive consumer. The naive consumer initially does not buy the biggest package (x_1^S) but the medium one

x_0^* – believing that he will stick to this quantity. But later, when in the “hot state" he buys a small top-up $(x_0^* - x_1^*)$. This results in inefficiently high shopping costs.

1.3.3 Monopolist

The monopolist maximizes his profits, $t_0 + t_1 - c\,(x_0 + x_1)$ subject to the respective participation constraints. For the sophisticated consumer, optimization problem (1.7) is unchanged, except for the reservation utility of self 0. Thus, as in the competitive market, self 0 receives either full or partial commitment. What changes is the transfer, which is given by $t_0 = u(x_0) - u(x_1^*) + v(x_1^*) \equiv u(x_0) - \bar{u}^S$, $x_0 \in \{x_0^*, x_0^S\}$.[9]

Consider next the naive consumer. Unlike firms in the competitive market, the monopolist can commit to sell to only one of the selves. He maximizes profits subject to the participation constraint of self 0 and/or self 1. As we show in the appendix, the monopolist never finds it optimal to sell a positive quantity to both selves. That is, he either caters to self 0 (by selling x_0^*), or self 1 (by selling x_1^*). Whether it is optimal to cater to the interest of self 0 or self 1 depends on whether $v(x_1^*) - c\,x_1^* - k_1$ is smaller or larger than $u(x_0^*) - c\,x_0^*$. While the sophisticated self 0 always does all the shopping, the naive self 0 never buys x_1^* and thus self 1 does the shopping in case it is optimal for the monopolist to cater self 1. This is due to the naïveté of the consumer. Self 0 does not anticipate that self 1 will buy x_1^*. Thus, at time 0, he will not buy x_1^* even if that could safe him the higher shopping costs k_1 at time 1. Thus, in this case, naïveté results in inefficiently high shopping costs.

Proposition 1.2.

1. *If $v(x_1^*) - c\,x_1^* - k_1 \geq u(x_0^*) - c\,x_0^*$ the monopolist offers contract (x_1^*, t_1^*), with $t_1^* = v(x_1^*) - k_1$ to the naive consumer, who goes shopping in period 1. If $u(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$ the monopolist offers contract (x_0^*, t_0^*), $t_0^* = u(x_0^*)$ to the naive consumer, who goes shopping in period 0.*

[9] Consider the subgame following a rejection of self 0 in period 0. In this subgame, the monopolist maximizes his profits by offering quantity x_1^* at transfer $t_1 = v(x_1^*) - k_1$. Self 1 accepts such an offer. This yields utility $u(x_1^*) - v(x_1^*)$ to self 0. Thus, self 0 accepts contracts with $u(x_0) - t_0 \geq u(x_1^*) - v(x_1^*) \Leftrightarrow u(x_0) - u(x_1^*) + v(x_1^*) \geq t_0$ given that $v(x_0) \geq v(x_1^*) - c\,(x_1^* - x_0) - k_1$.

2. *If* $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$ *the monopolist offers contract* (x_0^*, t_0^S)*,* $t_0^S = u(x_0^*) - \bar{u}^S$ *to the sophisticated consumer. Otherwise he offers* (x_0^S, t_0^S)*,* $t_0^S = u(x_0^S) - \bar{u}^S$ *and* x_0^S *defined by* $v(x_0^S) - c\,x_0^S = v(x_1^*) - c\,x_1^* - k_1$*.*

Hence, the monopolist either offers one package size (x_0^*), which provides perfect commitment to both types of consumers. Such perfect commitment is possible if the shopping costs of self 1 are large. Or he offers a small and a large package (either the pair (x_0^*, x_1^*) or the pair (x_0^S, x_1^*) depending on whether or not perfect commitment for the sophisticated consumer is possible). Thus, in a monopolistic market, small packages are always commitment devices for sophisticated consumers.

The solution for the sophisticated consumer is identical in the monopolistic and competitive market. Only the distribution of rents differs. Thus, the social surplus is the same. What about the naive consumer? He receives full commitment more often in the competitive market than in the monopolistic market. If, however, $v(x_1^*) - c\,x_1^* - k_1 \geq v(x_0^*) - c\,x_0^*$, so that not only self 0, but also self 1 goes shopping in the competitive market, then the total quantity consumed is equal in the monopolistic and the competitive market, but total shopping costs are higher in the competitive market. Thus, the social surplus is lower in the competitive than in the monopolistic market as the naive consumer goes shopping too often in the competitive market.[10]

1.4 Extensions and robustness

In the following section we wish to explore what happens if the model is extended in different ways. First, we discuss robustness to relaxing the assumptions of an observable degree of sophistication, on-the-spot-production and to allowing for partial naïveté. We do so for the competitive market.[11] Then we will look at the number of firms and degree

[10] Comparing the sophisticated and the naive consumer in the monopolistic market shows that the social surplus generated from the sophisticated consumer is larger than the one for the naive self 0. The monopolist offers full commitment to the sophisticated self 0 whenever $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$ and partial commitment otherwise. In contrast, the naive individual receives full commitment whenever $u(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$, i.e., as $v(x_0^*) - c\,x_0^* > u(x_0^*) - c\,x_0^*$, less often than the sophisticated individual.

[11] For the monopolistic market, we discuss the robustness of the results to the assumption that firms know the degree of sophistication and on-the-spot-production in the appendix. The arguments for

of competition.

1.4.1 Relaxing the assumption of observable degree of sophistication

The assumption that firms in the competitive market can distinguish between the naive and the sophisticated consumer is not crucial. Suppose three packages are offered. The commitment quantity x_0^S at the given transfer is not attractive for the naive consumer as he does not demand commitment. And $x_1^* - x_0^*$ and x_0^* are not attractive for the sophisticated consumer as they would induce overconsumption. If one package size (x_0^*) is offered, then it is offered at the same transfer to the sophisticated and naive consumer.

1.4.2 Relaxing the assumption of on-the-spot-production

The assumption that firms produce on the spot and can thus react to a period-0-deviation in period 1 neither drives the results for the sophisticated consumer if perfect commitment is feasible, nor for the naive consumer. In all these cases, no firm had an incentive to deviate – independent of what will happen in period 1. If however the sophisticated consumer receives only partial commitment, i.e., x_0^S, the assumption matters. Self 0 would prefer to buy a lower quantity than x_0^S – but only if he knew that self 1 would not go shopping. In the model with on-the-spot-production, firms can, in reaction to a period-0-deviation of some firm to say some lower quantity than x_0^S, offer self 1 a quantity that makes a shopping trip attractive. Anticipating this, self 0 would not buy the lower quantity and a deviation would not pay off.

If however firms could not react in period 1, the threat to sell self 1 some quantity might not be credible anymore. And if self 0 knew that self 1 would not go shopping, then he would buy the lower quantity from the deviating firm. But even if firms could not react to the deviation by tailoring a package for self 1, self 1 might still go shopping and buy one of the other available packages ($x_1^* - x_0^*, x_0^*$, or x_0^S). Anticipating this self 0 would not buy from the deviating firm. So if there exists no $x_0^* \leq x_0 < x_0^S$, such that $v(x_0) > \max\{v(x_0 + x_0^*) - c\, x_0^*, v(x_0 + x_0^S) - c\, x_0^S, v(x_0 + x_1^* - x_0^*) - c\,(x_1^* - x_0^*)\} - k_1$, offering

the former case are more involved, but the main insights are robust.

package x_0^S is still a competitive equilibrium. If this condition fails to hold, a firm has an incentive to deviate, and no competitive equilibrium exists.

1.4.3 Partial naïveté

In the main model we assume that there are two types of consumers: perfectly naive, and perfectly sophisticated consumers, implying $\hat{v}(x) \in \{u(x), v(x)\}$. In this section, we wish to explore if our main findings are robust to allow for partial naïveté. To do this we will assume $\hat{v}(x)$ can be anything between $u(x)$ and $v(x)$. The details can be found in the appendix. The implication is that consumers are aware of their self-control problem, but all except the perfectly sophisticated consumers underestimate the severity. This means that consumers are willing to pay for commitment, as we have shown in the baseline model. As in the baseline model, if x_0^* is sufficient to prevent perfectly naive consumers from shopping at time 1, all consumers receive full commitment. This is an intuitive extension of the baseline model. If x_0^* is sufficient to prevent the fully naive consumer from shopping at time 1, then the quantity also prevents partially naive consumers from shopping. If, on the other hand, x_0^* is not sufficiently big, then partially naive consumers do not buy a sufficiently large quantity to avoid shopping at time 1. Therefore, the consumers ends up consuming x_1^* but having paid more than $v(x_1^*) - k_1$, making them worse off than even the perfectly naive consumers. In fact, the closer to perfectly sophisticated, the greater the exploitation.

Thus, the result of (partially) naive consumers' too frequent shopping goes through. Yet, if there is a continuum of possible degrees of naïveté, then firms offer a continuum of quantities. Introducing a (plausible) fixed cost of offering a quantity-transfer pair into the model restores however a finite number of packages.[12]

[12] If we maintain the assumption of perfect competition, the number of offers will be such that the the marginal consumer of a quantity will be indifferent between buying and not entering the market, that is, even with perfect competition, the marginal consumer will not be left with any rents. In this setting, there will exist some level of sophistication such that all consumers at least as sophisticated as the cutoff will receive at least partial commitment, that is, the marginally naive will no longer be the most exploited.

1.4.4 Circle model

We now look at what happens if we change the competitive setting and allow for imperfect competition and market entry. To do this we will use the circular city model first presented by Salop (1979). In this setting, consumers are distributed uniformly along a circle where firms, when entering, choose a locations. This setting alters the nature of consumers' shopping costs in that the shopping costs now not only depend on the parameter k_τ but also on the distance θ to a firm. We assume that firms occur some fixed cost of entry into the market. Details can be found in the appendix.

Since we now have a fixed cost of entry, each firm must in either period serve a fraction of consumers and not just one individual for entry to be profitable. Hence if a type of consumers constitutes too small a fraction of consumers, it is no longer be possible to target them. If we abstract from this case, the results of the main model goes through: sophisticated consumers always receive at least partial commitment, and naive consumers are sometimes offered commitment, but only if x_0^* is sufficient to prevent shopping at time 1. The new aspect in this model is entry. Generally the fixed cost of entry along with free entry will drive prices up even if no firm makes positive profits, and in equilibrium entry is above the socially optimal level. This is not specific to our setting, but is a general result from the circle model.
What is specific about this setting is that each firm compete over consumers in either of the two periods. That is, firms operate on two separate markets: the one at time 0 and the one at time 1. This can be seen as reflecting the difference between offering a product in standard supermarkets, or in late night convenience stores.

1.5 Conclusion

The paper considers the shopping and consumption decision of an individual who faces a self-control problem. It asks how firms react to consumers' self-control problems, and possibly exploit them, by offering different package sizes. In a monopolistic market, small packages are commitment devices. In a competitive market, small packages may not be commitment devices, but, quite to the contrary, can serve to exploit naive consumers, who go (unexpectedly) frequently shopping. Further, the paper shows that while the sophisticated consumer receives the same commitment in the monopolistic and competitive market, the social surplus from the naive consumer can be lower in the competitive market because the naive consumer goes shopping more often in competitive markets.

Our model is consistent with recent results in the literature that suggest positive welfare effects of restricting sales times or locations of sinful products (see e.g., Beshears et al., 2006; Hinnosaar, 2012). Such policies increase k_1 and therefore make commitment easier to achieve, i.e., firms are less likely to offer products that cater the interests of the short-run self. Care should however be taken when restricting package sizes or subsidizing small packages. Depending on the market environment and the extend of the consumers' self-control problem, a large package can either be a commitment device or can serve the interests of the short-run self. Thus, a careful market analysis would be needed to decide which one is the case.

Chapter 2

The limited Potential of Partial-Commitment

2.1 Introduction

An extensively studied behavioral bias is present-bias: People's tendency to value the now disproportionately over the future. This bias implies that agents have time-inconsistent preferences, that is, agents may experience preference reversal. To overcome the issue of preference reversal, agents may take up commitment to (partially) commit themselves to a certain action. However, commitment is empirically rare[1] which poses the question: why do present-biased agents not demand commitment?

Commitment devices are devices or contracts that seek to align an agent's long-run and short-run preferences[2] or restrict the short-run agent's choice set. In the literature, agents who are aware of their present-bias are referred to as sophisticates and agents who are not, as naifs. If agents are naive, they are unaware that they have a self-control issue and therefore see no benefits to a commitment device. The natural conclusion would be that people are often naive while rarely sophisticated. In this paper, I will argue that even if agents are perfectly informed about their present-bias, they are often unlikely to take up commitment.

In my model, agents are faced with a task which has immediate (stochastic) cost and (fixed) delayed benefits. Agents know the cost distribution, but cannot contract on

[1] See Bryan et al. (2010) for an overview.

[2] This paper will only deal with quasi-hyperbolic discounting, or the so-called, $\beta - \delta$-model, and not general hyperbolic discounting or any other types of time-inconsistent preferences.

cost realizations. Therefore, the best the agents can do is to impose a punishment on themselves if they fail to complete the task. However, if there are cost realizations for which the cost outweighs the benefit even for a time-consistent agent, the agent is facing a tradeoff. Either ensure the task is only completed when the benefit is higher than the cost, but then incur the punishment with positive probability, or complete the task no matter the cost,even when the cost outweighs the benefit with positive probability.

Therefore, I focus on *partial* commitment, which here is defined as commitment that only ensures the task will be completed for some and not all cost realizations. The alternatives to partial commitment are no or full commitment, but I also introduce the term *perfect* commitment, which I define as commitment that perfectly aligns an agent's long-run and short-run preferences and thereby ensures that a time-inconsistent agent behaves as a time-consistent agent would. I also allow for agents to have a distribution of beliefs over their present bias. This modeling choice allows me to explore not only the importance of overconfidence but also underconfidence and precision in beliefs.

To my knowledge, this paper is the first to look at commitment devices with a socially wasteful component and this leads to even perfectly sophisticated agents optimally choosing non-perfect commitment devices. This is an important result since it implies that fully sophisticated agents may behave differently from time-consistent agents, even if they have taken up commitment. Therefore, seeing different behavior is not necessarily proof that an agent is partially naive and has picked the wrong commitment. I also find that if agents can make complete contracts, that is, contract on realized costs, commitment devices are attractive if agents know their present bias, but if costs are not contractable, agents only have the choice of taking up simple commitment contracts where agents have to pay a fee if they fail to complete a task, commitment uptake becomes costly.

To fix ideas imagine a simple everyday task that requires immediate effort with delayed benefits, for instance, going for a run. I start with perfectly sophisticated agents, and then analyze the implication of allowing for partial naïveté, which is often proposed as an explanation for little demand for commitment (For an overview, see Bryan et al., 2010 for an overview). One of the few suppliers of commitment devices is www.StickK.com — which is a company founded by behavioral economists. On this webpage you set a goal, (potentially) choose a punishment or add additional pressure from getting a referee and/or add friends.

Overall, my contribution is threefold: I propose a simple explanation for the limited commitment uptake – that commitment can be socially wasteful and therefore often undesirable, that commitment uptake is non-monotonically related to the degree of present-bias, and that agents have limited value of knowing their exact present-bias – offering a potentially explanation for why naïveté is empirically prevalent.

2.2 Literature

Although the literature on self-control problems goes back to Strotz (1955) and Phelps and Pollak (1968), and the models of present biased were formalized in the late nineties with Laibson (1997) and O'Donoghue and Rabin (1999), present-bias is still an intensely studied topic to this day.[3]

The papers in this literature that are conceptually most closely related to this paper are these by DellaVigna and Malmendier (2004), and Laibson (2015). Both of these examine commitment uptake and both papers use settings with stochastic costs.

The paper by DellaVigna and Malmendier (2004) examines the interaction between present-biased agents and profit-maximizing firms. The paper is not explicitly modeling commitment uptake, but derives the optimal contracts for firms selling either investment or leisure goods to present-biased consumers. The type of tasks my paper is describing is exactly investment goods, meaning goods or tasks where there is an immediate cost and a delayed benefit. Therefore, the self-control problem and payoff structure is the same, abstracting from the potential commitment. In their paper, firms price an investment good below marginal costs such that preferences of time-inconsistent agents are aligned with those of a time-consistent agent, but for high costs realizations the agent can choose not to buy and will not incur any punishment. This implies that the contracts firms offer to sophisticated consumers will make them fully overcome their present-bias, while partially naive consumers will be exploited. In my paper, being perfectly sophisticated is

[3] For an overview see Bryan et al. (2010).

not enough to overcome one's present-bias since commitment in my model has a wasteful component — the fine the agent has to pay when not completing the task.

Laibson (2015) also seeks to explain why little commitment uptake is observed and he also shows that commitment should be rare in many real-world settings. There are, however, important differences in the settings. He is looking at a task that *has* to be completed, and where the choice is when to complete it. He assumes a flow loss from not completing task in any given period along with stochastic costs of completing the task (exactly as in my model and as in DellaVigna and Malmendier's paper), and this is what drives the intra-personal conflict in his setting.

Laibson (2015) defines commitment as limiting of choice sets (In this particular case, by setting binding deadlines). Though this is often how commitment is implemented in experiments, I would argue that outside of the lab, this is a rarely observed type of commitment. He finds that agents either allow themselves full flexibility to pick the time to complete the task, or impose a binding deadline in the very first period. In essence, this is a choice between no commitment and full commitment. In my setting, partial commitment is an option. It is not straight-forward to compare results across the two frameworks, but in Laibson's (2015) paper, he finds that increasing the variance of costs decreases demand for commitment. I find that the demand for partial commitment increases with the variance in costs while having no effect on demand for full commitment.

In both of these papers, among others (Eliaz and Spiegler, 2006, O'Donoghue and Rabin, 2001), partial naïveté is modeled as either a point belief about the degree of present-bias, or as a simple probability an agent attaches to being time-consistent. My model allows for a very flexible definition of partial naïveté where agents have a distribution of beliefs about their present bias rather than an overconfident point-belief about their type. I find that underconfidence rather than overconfidence is certain to decrease demand for commitment.

In section 2.5 I use a simple finite horizon adaptation of the model by Ali (2011) to give an example of the tradeoff between commitment and learning. This is also the tradeoff of interest in his paper, but there are some important differences: as in the paper by

Laibson (2015), Ali is mainly interested in restricting choice sets[4]. He does also allow for so-called "nudges" that are similar to the commitment devices I propose, with the important distinction that these commitment devices have no payoff consequences for the "planner". I will argue that the value of learning, or of experimenting, is lower in my setting, and, unlike in his paper, non-informative under-commitment can happen persistently (though this follows from a difference in assumptions). My setting therefore also allows for persistent overconfidence.

Generally, the topic of persistency of biased beliefs have received increasing attention. Christensen and Murooka (2018) explicitly looks at non-learning about present-bias in an imperfect-recall model without commitment devices. In this paper, non-learning can be sustained if agents are sufficiently optimistic about their future behavior and therefore perceive it of little importance to acquire information today. Carrillo and Mariotti (2000) show how strategic ignorance may be used as a self-discipline device, and ignorance or biased beliefs function as a commitment device to restrict consumption of vice goods. Bénabou and Tirole (2002) analyze the benefits of self-serving beliefs and characterize when overconfidence can persist, while Bénabou and Tirole (2004) use self-signaling about own willpower as a commitment device, which can lead to persistent under-confidence. I allow for biased beliefs about self-control in either direction, and can ensure persistency in either direction.

The structure of the remainder of the paper is as follows: In Section 2.3 I introduce the baseline model, in Section 2.4 I analyze the consequences of partial naïveté, in Section 2.5 I consider commitment uptake in a dynamic setting, in Section 2.6 I discuss implications and possible further directions, and in Section 2.7 I conclude.

2.3 Model

Consider a risk-neutral agent contemplating whether to go for a run. Running will have immediate costs $c \in [\underline{c}, \overline{c}]$ where c is distributed according to cdf $F(c)$, and delayed health benefits b. The agent is present-biased, such that the agent discounts future payoffs by

[4] Building on the modeling of temptation from Gul and Pesendorfer (2001)

$\beta\delta^t$. I assume that $\delta = 1$, and $\beta \in [0, 1]$. Further, I assume that $\beta b - \underline{c} < 0$[5] and $b \in [\underline{c}, \bar{c}]$. Together, these assumptions ensure that the agent will never go running, even for the lowest possible costs, absent a commitment device, even though running is worthwhile for some realizations of c, but not for all. Imagine, for instance, that the agent thinks it is important to exercise regularly, and would consider running a worthwhile way of doing so, but if there is a snowstorm the potential health benefits are not worth the cost of the run. I will throughout the paper mainly focus on the case where $b < E[c]$. This implies that the agent would prefer no commitment over full commitment.

In the following I characterize the agent's preferred commitment device for the case where costs are contactable. This will serve as a benchmark for the case where costs are not contactable.

2.3.1 Benchmark

Agents would like to go running whenever $b - c \geq 0$, therefore, an agent would like to sign a commitment contract where there is no consequence of not running whenever $c > b$, but where there is a punishment if the agent does not run whenever $b > c$. If agents know their present bias, that is, if agents have correct beliefs about their β, agents would want any contract where the agent would have to pay fee $k(c) \geq (1 - \beta)b$ whenever $c \leq b$ and fee $k(c) = 0$ whenever $c > b$.

Lemma 2.1. *The optimal contract for a perfectly sophisticated agent is any contract where:*

1. $k(c) \geq c - \underline{\beta} b \quad \forall c \leq b$.
2. $k(c) = 0 \quad \forall c > b$.

There are three things to note from this. First, an agent never actually pays anything but only incur the cost c, whenever $c < b$. Therefore, agents behave as if they were time-consistent agents and the optimal commitment contract offers perfect commitment. Second, this implies that the degree of present bias has no utility implications, and

[5] This assumption can be relaxed and is only crucial in a repeated setting.

therefore that demanding commitment is independent of an agent's present bias[6]. Last, increasing the variance of costs, does not necessarily make the agent worse off. Under the assumption that $b < E[c]$, a mean preserving spread of $F(c)$ implies that the task is worthwhile with a higher probability, and that the expected gain, $b - E[c|c \leq b]$, is higher.

If β is not known, but the agent has beliefs about the distribution, $\hat{G}(\cdot)$, of β, where the domain of β is bounded and the lower bound $\underline{\beta} \geq 0$, then the agent is indifferent between any contract with a fee of $k(c) \geq (1 - \underline{\beta})b$ whenever $c \leq b$ and a fee of $k(c) = 0$ whenever $c > b$.

Lemma 2.2. *Given any beliefs $\hat{G}(\beta)$ about the distribution of β, the any commitment contract with the following fee structure $k(c)$ will implement the first best:*

1. $k(c) \geq c - \underline{\beta}b \quad \forall c \leq b$.
2. $k(c) = 0 \quad \forall c > b$.

Condition 1 and 2 are as derived above, and ensure that the commitment contract perfectly aligns the agent's long-run and short-run preferences. If an agent accepts such a contract (and is correct in believing that β is not smaller than $\underline{\beta}$) then the agent never actually incurs the fee k, and completes the task exactly when the patient long-run self would wish to have the task completed. The valuation of such a contract is $Pr(c \leq b)(b - E[c|c \leq b]) = F(b)(b - E[c|c \leq b])$ assuming the agent would never complete the task absent a commitment device. Since the value of the task is independent of β, and under the assumption that the task would never be completed if the agent did not choose some level of commitment, the value of commitment is not directly dependent on β. But since the cutoff for $k(c)$ depends on β, the degree of time-inconsistency and the uncertainty of beliefs over β matter.

Constructing a contract of the described form relies on c being observable and verifiable. Even if that is the case for most realizations of c or most states of the world (some

[6] Given that $\beta < 1$.

probability of not being observable is sufficient for these contracts to break down). The extreme situation is one where c is completely unobservable to a contractor.

Returning to the example of going for a run: Assume the agent considers it optimal to go running every Tuesday, unless there is a blizzard or the agent has a headache. A blizzard would be possible to contract on, and since weather conditions are easy to verify, it should be feasible to make the contract conditional on the weather. Headaches, on the other hand, are trickier to contract on. Therefore, I consider the case where c is not contractable to characterize the agent's optimal commitment contract (if one exists) in this setting.

2.3.2 Second Best Contract

If the contract cannot be contingent on c, the only feasible commitment contracts involve a fixed $k > 0$ if the agent does not complete the task. In the case of the running example, k could be an actual monetary amount to be paid to a third party on Stickk, it could be a fee for not showing up, if the commitment device is to run with a running club, or it could be a psychological cost, or reputation cost, if the agent told friends or colleagues that they would go running, and then failed to. I abstract from any costs of setting up commitment. That means, it does not cost the agent anything to make a contract on Stickk, to join a running club, or to tell friends about their plans to go running.

In the analysis when c is not contactable, the actual value of β becomes crucial along with the beliefs about β. In the following I start out by identifying the optimal contract for a fully sophisticated agent. Thereafter I examine the implications of introducing uncertainty over β either as biased or unbiased beliefs.

Perfectly sophisticated agents with non-contractable costs

In this section I characterize the optimal commitment contracts for a fully sophisticated agent when contracts cannot be contingent on c, where c is distributed according to cdf $F(c)$ with pdf $f(c)$. If an agent chooses a fee k, the agent has the payoff $\beta b - c$ when completing the task, and $-k$ if not completing the task. Hence, the agent completes the task whenever $\beta b - c > -k$. The agent does not know the cost realization c when choosing a k, but only the distribution of costs. The agent can freely choose any fee, k, for a commitment device and because the costs are bounded it is always possible to construct a commitment device (i.e. choose a k) such that the present-biased agent always completes the task i.e. full commitment is available.

I am interested in the case where the agent prefers no commitment to full commitment, that is $0 > b - E[c]$. Otherwise, if $0 < b - E[c]$ some level of (potentially full) commitment is always optimal. Therefore, under the assumption that $0 > b - E[c]$ if an agent demands commitment it is always an intermediate level of commitment. The payoff from a commitment device with fee k is given by:

$$U(k) = F(\beta b + k)(b - E\left[c|c \leq \beta b + k\right]) - k(1 - F(\beta b + k)).$$

The following proposition gives necessary and sufficient conditions for a positive commitment level k being optimal for some degree β of present-bias:

Proposition 2.1. *Assuming c is continuously distributed according to cdf $F(c)$ where $E[c] > b$. Then, if $\exists\, c \in (\underline{c}, \overline{c})$ s.t.:*

$$b > \frac{1 - F(c')}{f(c')} > -\frac{f(c')}{f'(c')} > 0, \tag{2.1}$$

then $\exists\, \beta'$ s.t. intermediate commitment $k = c' - \beta' b$ is optimal. For a given β', non-zero commitment is optimal if $\exists k' > 0$ s.t.:

$$b(1 - \beta) = \frac{1 - F(\beta' b + k')}{f(\beta' b + k')} > -\frac{f(\beta' b + k')}{f'(\beta' b + k')} > 0. \tag{2.2}$$

For details, see Appendix B.0.1. The conditions on the distribution of c require that the Mill's ratio of $F(c)$ is (locally) increasing, and that the pdf is (locally) decreasing. This is not the case for several important classes of distribution functions including Normal distributions and uniform distributions, though it is the case for Pareto distributions and inverse logistic distributions.

Therefore, whenever $E[c] > b$, a non-zero level of commitment is rarely optimal. However, if $E[c] < b$, by the same argument, full commitment is likely optimal, even though there are realizations of c where even a time-consistent agent does not want to complete the task.

Intuitively, when increasing the fee, k, the agent is facing a tradeoff between decreasing the likelihood of having to pay k, but paying a higher k, along with increasing the probability of completing the task but decreasing the expected gain of a completed task. Whether these tradeoffs lead to the agent choosing commitment that leads to over- or under-completion of the task relative to a time-consistent agent depends on the distribution of costs. It is possible that the agent prefers a commitment device k such that the time-inconsistent agent completes the task exactly when a time-consistent agent also completes the task. This is the case only if $k = b(1 - \beta)$, or:

$$b(1-\beta) = \frac{1-F(b)}{f(b)} > -\frac{f(b)}{f'(b)} > 0.$$

Since the right-hand side does not depend on β, for each possible cost distribution $f(c)$ where an interior optimum exists, there is full alignment of long-run and short-run preferences for at most one unique β. Therefore, perfect commitment is not a general result, but rather the exception. If $0 < -\frac{f(b)}{f'(b)} < \frac{1-F(b)}{f(b)} \leq b$, then there exist a β for which perfect commitment is optimal.

Corollary 2.1. *If agents face stochastic costs that cannot be contracted upon, then:*

1. *The optimal commitment contract does not ensure that sophisticates behave as time-consistent agents.*
2. *Perfectly sophisticated consumers may optimally choose no commitment.*
3. *Observing an agent behaving differently from a time-consistent agent is not proof that the agent is naive.*

Whether a sophisticated time-inconsistent agent is more or less likely to complete the task than a time-consistent agent depends on the parameters of the model. The important finding, however, is that it is not proof of an agent being naive, that an agent does not complete the task for a cost realization c for which a time-consistent agent would. It could be naïveté, but it could also be that perfect commitment is not optimal. This finding is in contrast to the benchmark case where perfect commitment is always optimal and also in contrast to the findings by DellaVigna and Malmendier (2004) where agents only demand non-perfect commitment if they are naive.

So far, I have considered continuous distributions of c where non-zero commitment is only optimal for a limited set of distribution functions and parameter values. In order to build intuition, I, now, consider simple discrete cost distributions to examine under which conditions non-zero commitment is optimal.

For c distributed according to the discrete cdf $F(c)$, with pdf $f(c)$, existence of non-zero partial commitment is ensured by the following proposition:

Lemma 2.3. *If c is distributed according to the discrete pdf $f(c)$, a perfectly sophisticated agent demands non-zero commitment if:*

$$\max_{c_i} U(k = c_i - \beta b) = \max_{c_i} \beta b - c_i + \sum_{c=\underline{c}}^{c_i} f(c)\left(b(1-\beta) + c_i - c\right) > 0, \tag{2.3}$$

where $c_i \in supp(f) = \{\underline{c} = c_1, c_2, \ldots, \overline{c} = c_n\}$, and $c_1 < c_2 < \ldots < c_n$.

Details can be found in Appendix B.0.2, but the intuition is the following: the agent prefers paying the lowest possible fee which ensures a certain action. Thus, any k an agent considers satisfies $k = c_i - \beta b$ for some $c_i \in \text{supp}(f)$. If no such k gives the agent a positive payoff in expectation, the agent prefers no commitment.

2.3.3 Example with two cost levels

In this section, using the result from Lemma 2.3, I provide a simple example for existence of optimal partial commitment for sophisticated agents. Furthermore, through a comparative statics I address the implications of changes to the underlying cost distribution.

If there are only two possible costs $c \in \{\underline{c}, \overline{c}\}$, where $\underline{c} < b < \overline{c}$, with probability p of $\underline{c}$ occurring, partial commitment implies that the time-inconsistent agent completes the task if $c = \underline{c}$ but not if $c = \overline{c}$.

Coming back to the running example, this is the case if agents would run whenever they do not have a headache, but stay at home otherwise. Therefore, the preferred intermediate commitment device makes an agent indifferent between running and not running when the agent does not have a headache. The cost could, for instance, be an amount of money to pay, or a number of people that the agent will have a reputation loss from.

Proposition 2.2. *If no commitment is preferred over full commitment, a perfectly sophisticated agent prefers intermediate commitment if:*

$$\beta \geq \frac{\underline{c} - b \cdot p}{b - b \cdot p}.$$

If full commitment is preferred over no commitment, a perfectly sophisticated agent prefers intermediate commitment if:

$$\beta \geq 1 - \frac{\overline{c} - \underline{c}}{b}.$$

If the agent chooses such a commitment device, the present biased agent's preferences will be aligned with those of a time-consistent agent. That is, any intermediate commitment devices in this setting ensures perfect commitment. This result is driven by the two-cost distribution, rather than being a general result.

In both cases, the right-hand side is strictly less than $\underline{c}/b(> \beta)$, which means that for this particular simple example, there is always a cutoff for β, such that every agent with a less severe present bias than the cutoff will choose intermediate commitment. See Appendix B.0.3 for details.

It is even possible that the condition will be satisfied for all positive β, meaning that even the most severe degree of present bias ($\beta = 0$) can be overcome such that all perfectly sophisticated agents behave as time-consistent agents. In the running example, this implies that anyone can find a commitment device that would make them run with certainty if they do not have a headache, but never otherwise.

The main points provided by this example are summarized in the following:

Corollary 2.2. *In a setting with two possible cost levels:*

1. *Demand for perfect commitment is decreasing in present-bias.*
2. *If no commitment is preferred over full commitment only the least present biased agents will demand commitment.*
3. *If full commitment is preferred over no commitment, every agent demands commitment but only the least present-biased behave as time-consistent agents.*
4. *An increase in the spread of costs increases the demand for perfect commitment.*

To understand the intuition, consider the running example: Preferring full commitment over no commitment occurs when the high cost, having a headache, is rare. In this case, all agents agree that running in the unlikely event of having a headache is preferred over never running. But for partial commitment, the different types disagree: if the agent is severely present-biased the fee, k, necessary to make the agent run is high, even when the agent does not have a headache. Therefore, the cost the agent will incur from not running with a headache outweighs the cost from running with a headache. Therefore, the severely present-biased will demand stronger commitment than the less present-biased.

If no commitment is preferred over full commitment, the least present-biased agents still demand perfect commitment but in this case, the most present-biased agents choose no commitment. The argument is similar. Intermediate commitment becomes too expensive for the severely present-biased because the fee necessary to ensure any running will be so high that paying it whenever the agent has a headache outweighs the benefit from running when healthy.

Overall, demand for commitment is ambiguously related to the agent's degree of present bias. It is monotonically related to present bias given a specific cost function, but not across cost functions. Additionally, the closer to time-consistent an agent is, the likelier it is that the agent can fully overcome the self-control issue and behave like a time-consistent agent.

Further, increasing the spread of c, keeping $E[c]$ fixed, increases the likelihood that the agent prefers the intermediate level of commitment over any of the extremes. Implying that a higher variance of costs benefits the agents unambiguously, at least weakly. The intuition is as follows: In this simple case, increasing the spread means decreasing $\underline{c}$. A lower $\underline{c}$ increases the net gain from completing the task when the agent is healthy, $b - \underline{c}$, holding fixed β. Therefore, a lower fee is needed to ensure the agent goes running when healthy, which means that the fee to be paid when the agent has a headache is lower.

An example with only two cost levels has its limitations. Especially since perfect commitment is the only partial commitment available. Three cost levels allows for characterization of the existence of over- and under-commitment. I have also examined this case and the main findings are the following:

Proposition 2.3. *In a setting with three possible cost levels:*

1. *A sophisticated agent may optimally commit to completing the task more or less often than a time-consistent agent.*
2. *A larger self-control problem ambiguously affects the demand for commitment.*
3. *Increasing the spread of costs increases the likelihood of a sophisticated agent demanding perfect commitment and demanding partial commitment overall.*

Details can be found in Appendix B.0.4.

So far, the agents have been perfectly sophisticated, which is a strong assumption. In the following, I examine the implications of relaxing this assumption.

2.4 Partial naïveté

In this section I examine the importance of the most prominent assumption in the present-bias literature: naïveté. Often naïveté is modeled as a point belief, $\hat{\beta}$, about β where $\hat{\beta} > \beta$. The advantage of this assumption is its simplicity while still explaining many core findings. However, I allow for a more flexible definition of partial naïveté: agents can have any distribution of beliefs, $\hat{G}(\beta)$, over β.

Partial naïveté in this context therefore does not necessarily imply over-confidence, but allows for both under- and over-confidence in expectation. Thereby, this definition also allows for analysis of the implication of uncertainty itself. I use the term *partial naïveté* to refer to any non-degenerate beliefs about β.[7]

The valuation of an intermediate commitment device k is:

$$U(k) = Pr(\beta b - c \geq -k)(b - E[c|\beta b - c \geq -k]) - k \cdot Pr(\beta b - c \geq -k),$$

where the only difference from the perfect sophistication case is that, in this case, the agents have beliefs about how β and c are distributed.

Define $z = c - \beta b$ and let $F_z(x)$ and $f_z(x)$ be the cdf and pdf of z, then the agent's expected value of an intermediate commitment device k is:

$$U(k) = F_z(k)b - \int_{\underline{c}}^{\overline{c}} cf(c)\left(1 - \hat{G}\left(\frac{c-k}{b}\right)\right) dc - k(1 - F_z(k)).$$

Which leads to the following optimality condition:

$$(k+b)f_z(k) = \frac{1}{b}\int_{\underline{c}}^{\overline{c}} cf(c)\hat{g}\left(\frac{c-k}{b}\right) dc + 1 - F_z(k).$$

The left hand side is the marginal gain from increasing k, that is, the agent is more likely to complete the task and therefore to receive b, while being less likely to incur the fee k. On the right hand side are the costs of increasing k. The agent incurs the marginally higher k whenever the agent does not complete the task, or with probability $1 - F_z(k)$. Additionally, the expected cost is now higher, conditional on the task being completed. Overall, the agent's problem is very similar to that of known β, only with added uncertainty. To develop intuition for the implications of such uncertainty and the tradeoffs the agent is facing, I go through a simple example.[8]

[7] I therefore abstract from the case where agents are certain about their type, but have wrong beliefs.

[8] Details can be found in Appendix B.0.5.

2.4.1 Example with unknown present-bias and three cost levels

Consider a setting where β is uniformly distributed, with $\beta \in \left[\underline{\beta}, \overline{\beta}\right]$, $c \in \{c_1, c_2, c_3\}$ with $c_1 < c_2 < c_3$ and the probabilities of the three cost realizations are p_1, p_2, p_3, respectively.

Assume that $c_1 < b < c_2$, $b < E[c]$ and $c_2 - \underline{\beta}b < c_3 - \overline{\beta}$. These assumptions ensure that a time-consistent agent only wants to complete the task for the lowest cost, c_1, that the task will never be completed without commitment, and that there exist a k such that the task will be completed with certainty for c_1 and c_2, but never for c_3.

The agent's utility from choosing a commitment device with a fee, k, is given by:

$$\begin{aligned} U(k) = & p_1\left((b-c_1)Pr(\beta b - c_1 \geq -k) - k \cdot Pr(\beta b - c_1 \geq -k)\right) \\ & + p_2\left((b-c_2)Pr(\beta b - c_2 \geq -k) - k \cdot Pr(\beta b - c_2 \geq -k)\right) - k \cdot p_3. \end{aligned}$$

If there exists a $k > 0$ such that this is strictly positive then the agent demands commitment. If the agent's preferred k ensures that the task will be completed for c_1 and never otherwise, then the optimal commitment is perfect. Because the agent has uniform beliefs over β, any optimum is a corner solution[9], but because of the discrete nature of the cost distribution function, there are not just the two extremes of full or no commitment; instead there are additional kinks.

The following proposition gives sufficient conditions for an agent demanding perfect commitment:

Proposition 2.4. *A partially naive agent will choose perfect commitment if the following holds:*

$$\underline{\beta} \geq \max\left\{1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}, \frac{c_1 - b \cdot p_1}{b - b \cdot p_1}\right\},$$

under the assumptions that $\overline{\beta}b - c_1 < 0$ *and* $b(\overline{\beta} - \underline{\beta}) < \min\{c_3 - c_2, c_2 - c_1\}$.

The last assumption ensures that it is possible to find a commitment device such that

[9] The joint distribution function of β and c is piecewise uniform. Therefore, the derivative of the joint pdf is non-negative and violates the second order condition for utility maximization.

the agent completes the task with certainty for $c = c_1$ but never for $c = c_2$, or with certainty for $c \in \{c_1, c_2\}$ but never for $c = c_2$.[10] Therefore, there are two interior kinks: the smallest k that makes the agent complete the task with (perceived) certainty for c_1 and otherwise not, and the smallest k that makes the agent complete the task for c_1 and c_2 but not for c_3. Under these assumptions, the agent can be certain of the action taken given any realization of costs and the agent therefore has to choose whether to never complete, complete only if $c = c_1$, complete if $c \in \{c_1, c_2\}$ or always complete the task. In Appendix B.0.3 I analyze the case where the assumption does not hold.

Overall there are two minimum bounds for $\underline{\beta}$, implying that under the simplifying assumptions, an agent is less likely to take up a perfect commitment device the bigger a present bias issue the agent believes to have with positive probability. In this example, neither the spread, nor the upper bound of the support of the beliefs directly affect the agent's commitment uptake. This implies that uncertainty about own self control issues itself is not a deterrent for commitment device uptake, and neither is overoptimism or biased naïveté.

However, perfect commitment is not the only option. It is possible the agent prefers a k such that the task will be completed even for c_2. The following proposition provides sufficient conditions for a partially naive agent choosing stronger than perfect commitment.

Proposition 2.5. *A partially naive agent will choose a commitment device that is stronger than perfect commitment if:*

$$\underline{\beta} \in \left(\frac{c_1}{p_3 b} + \frac{c_2 - c_1 - b}{b} \cdot \frac{p_2 + p_3}{p_3}, 1 - \frac{c_2 - c_1}{b} \frac{p_2 + p_3}{p_2} \right),$$

under the assumptions that $\overline{\beta} b - c_1 < 0$ and $b(\overline{\beta} - \underline{\beta}) < \min\{c_3 - c_2, c_2 - c_1\}$.

In this case, the agent prefers the stronger commitment device even though it means completing the task when the cost is higher than the benefit. It is not clear how demand for commitment behaves with an increase in spread of β or when more or less overconfident agents would demand such commitment, but it is clear that there is a lower bound on $\underline{b}$

[10] Details can be found in Appendix B.0.6.

such that for β lower than this, an agent chooses no commitment. Therefore, increasing the spread while keeping the mean of β constant does decrease the agent's propensity to acquire commitment.

Proposition 2.6. *Under the assumption that an agent believes the task will not be completed without commitment, the example shows that:*

1. *There is no direct effect of the mean β on demand for commitment.*
2. *Demand for perfect commitment is increasing in the lower bound of beliefs on β.*
3. *There is no direct effect of uncertainty about β on demand for commitment.*
4. *A mean preserving increase weakly decreases demand for commitment.*

However, increasing the uncertainty over β eventually leads to violation of the assumption $b(\overline{\beta}-\underline{\beta}) < c_2 - c_1$. The details for this case can be found in Appendix B.0.6. I will, in the following, give brief summary of the main differences between this case and the former.

Lemma 2.4. *Without perfect separability, the following holds:*

1. *Perfect commitment is no longer possible in expectation.*
2. *The agent faces a tradeoff between ensuring task completion for low costs but risk completing for high costs, and prevent completing for high costs but but risk not completing for low costs.*

Overall, this example shows that being overconfident, either defined by $E[\beta]$ or $\overline{\beta}$ does not directly affect the demand for commitment. More important is the biggest self-control problem an agent attaches a positive probability to having. Being under-confident, or overestimating one's own self-control problem ma,y however, be a bigger hindrance for overcoming self-control problems than underestimating it. The intuition is again that the fee needed is too high for intermediate commitment to be worthwhile, and agents therefore gravitate towards either full or no commitment. In this example I imposed an upper bound on the agent's possible overconfidence. I assumed that $\overline{\beta}b - c_1 < 0$, which means that the task will not be completed even for the smallest present-bias that the agent perceives possible. This is a strong assumption, and if I relax this assumption there exist beliefs $\hat{g}(\beta)$ such that the agent prefers no commitment. In the limit, if the agent has degenerate beliefs such that the agent believes $\beta = 1$ with certainty, the agent will not demand any commitment.

Furthermore, a marginal underestimation of present bias leads the task to be completed too rarely, therefby incurring the fee too often, but also means that the fee is slightly lower. Partial naïveté in the form of underconfidence may make an agent not take up commitment at all, and since the optimal choice of commitment is sometimes solely determined by the lower bound of β, a small positive probability attached to a very low β might be detrimental, even if the agent is overconfident in expectation. This also means that there is a non-monotonic demand for commitment as a function of β.

Even though this is a very simple example, where the simplicity of the results depend on the assumption of uniform costs, this might still help explain why there appears to be persistence of naïveté. Because the optimal action for an agent only indirectly depends on the uncertainty about preferences, there is little incentive to form precise beliefs, and since only pessimistic beliefs are unambiguously bad for the agent, while optimistic beliefs might be completely innocuous, there is also little incentive to form unbiased beliefs.

To show the intuition for the limited incentive to learn, I give an example of a simple dynamic decision problem where an agent may optimally stay naive.

2.5 Example of Learning about Self-Control

So far I have considered an agent's choice of commitment in a static setting. But for many important real world tasks, like running, the agent can complete them multiple times. In a repeated setting where agents are not only considering whether to complete a task once, agents potentially value experimentation if this can lead to more precise beliefs about their present bias.

I go through a finite horizon adaptation of the setup from Ali (2011) but unlike in his framework it is not possible to directly restrict choice sets. In his paper, Ali does allow for part of the commitment to come from costs that are triggered by a failure to complete the task, but an important difference is that these costs are not directly payoff relevant for the planner. That is, if the agent picks a commitment device that does not provide full commitment, the agent does not take into account the expected cost from paying the fee with positive probability. In this paper, the core assumption is that the agent has to pay this fee if the task is not completed. I also allow for perfect recall, so that the agent can look back and see exactly which cost realizations along with which commitment devices have been successful in the past. Despite this, I find that agents may not learn. As in the one-shot case, the agent's preference over full commitment versus no commitment is independent of the agent's beliefs about their own present bias. Only the expected value of partial commitment depends on the agent's beliefs.

It follows directly from the analysis in Ali's (2011) paper that once the agent chooses a non-informative option, here no commitment or full commitment, the agent sticks to this option indefinitely. This implies that if no commitment is preferred over full commitment, the agent will either keep experimenting until the agent is perfectly informed about their type, or switch in finite time to no commitment and stay there indefinitely with no further updating of beliefs. Given my assumption that no commitment is uninformative, this also means that, unlike in Ali's paper, no commitment combined with limited learning can be a long-run outcome, although, in line with Ali's paper, this happens when the agent reaches a sufficiently pessimistic posterior.

2.5.1 Finite period example

In order to get an intuition for the dynamics at play. Consider a two period setting, $T = 2$, with two costs, $c \in \{\underline{c}, \overline{c}\}$, two possible degrees of present bias, $\beta \in \{\underline{\beta}, \overline{\beta}\}$, and with $\delta = 1$.

Going back to the running example. Imagine the cost are that either the agent has a headache, or the agent is in good health. The agent believes that they will not go running without some commitment, but are unsure about how much commitment would be needed. With probability g the agent believes they have a present bias parameter of $\underline{\beta}$, and with $1 - g$, $\overline{\beta} > \underline{\beta}$. The agent knows that the chance of not having a headache is p, and therefore there is a risk of $1-p$ of having a headache. If the agent were to take up perfect commitment, it would either be $\underline{k} = \underline{c} - \overline{\beta}b$ or $\overline{k} = \underline{c} - \underline{\beta}b$. I assume that $\overline{\beta}b - \overline{c} < \underline{\beta}b - \underline{c}$, such that the agent can ensure they will run when healthy without risking that the commitment is so strong that they will also run with a headache with strictly positive probability.

As in the previous sections, I assume that $b < E(c) = p\underline{c} + (1-p\overline{c})$ so that the agent would prefer never running over having to run every time they have a headache. If the agents were completely myopic, they would pick either no commitment, or the intermediate commitment device that maximizes their expected payoff in the current period. If no commitment is preferred over either of the two intermediate commitment levels then the myopic agent would choose the non-informative action of $k = 0$, or no commitment, in both periods.

But even though agents might not want any intermediate commitment in the short-run given their prior beliefs, they might want to experiment to achieve a better expected outcome over both periods. In this simple example, the only way to experiment is to pick a commitment device such that there will be different actions taken for the two possible levels of β, at least given some cost realization. A partially naive agent chooses costly experimentation in the first period if the following holds:

Proposition 2.7. *A partially naive agent chooses experimentation in a two-shot game if:*

$$\underline{\beta} < \frac{\underline{c} - b \cdot p}{b - b \cdot p},$$

and:

$$\overline{\beta} \in \left\{ \frac{\underline{c} - b(1-g)p - (b - \underline{c})p(1-g)}{b - b(1-g)p}, \frac{\underline{c} - b(1-g)p}{b - b(1-g)p} \right\}.$$

If the agent picks the higher level of commitment, $\overline{k}$, the agent always goes running if healthy, and never with a headache, therefore, this option does not provide any information about the agent's present bias. The lower level of commitment, $\underline{k}$, does potentially provide information. If $c = \overline{c}$ the agent never goes running, so here no additional information is provided, but if $c = \underline{c}$ the agent only goes running if the present-bias is small,

$\overline{\beta}$, and not with a bigger present-bias, $\underline{\beta}$. Therefore, if the agent prefers experimentation, the agent will pick $\underline{k}$.[11]

However, if $\underline{k}$ is also the optimal choice in a one-shot game, learning is a side-effect of the short-term optimal decision, and not a sign that the agent values experimentation. The upper bounds on β ensures that the agent would not choose commitment in a one-shot game, but would choose intermediate commitment if $\beta = \overline{\beta}$.

This means that if the agent only considered going running once, the perceived probability of having the less severe present bias, $\overline{\beta}$, is too low for intermediate commitment to be optimal, but if the agent knew $\beta = \overline{\beta}$ intermediate commitment is optimal. Therefore, the information has potential value to the agent.

If the agent does not experiment, the best option is choosing no commitment and having zero payoffs in both periods. The only option for experimentation is choosing $\overline{\beta}$ at $t = 1$ and then choosing either no commitment or intermediate commitment at $t = 2$, depending on the outcome from the first period. If the agent did go running at $t = 1$, the agent knows for sure that $\beta = \overline{\beta}$. Not completing the task if $c = \underline{c}$ implies the agent has $\beta = \underline{\beta}$ with certainty, while not completing the task for $c = \overline{c}$ provides no information. Experimenting leads the agent to incur an expected negative payoff in the first period, but with positive probability to incur a strictly positive payoff in the second period.

From this simple example are the following take-aways:

Corollary 2.3. *If no commitment is preferred over full commitment:*

- *A partially naive agent may stay naive.*
- *Propensity to experiment is decreasing in expected present-bias.*
- *Propensity to experiment is increasing in spread of present-bias.*

It is in itself not a surprising result that the agent will experiment under some conditions if the task is repeated.

[11] For details, see Appendix B.0.8

What is interesting is that choosing no commitment is likely the more present biased the agent believes to be, and that no commitment can lead to no learning. If the low level of intermediate commitment $\underline{k}$ is optimal even in a one-shot, the agent continues to choose this level of commitment indefinitely if $\beta = \overline{\beta}$, and the agent becomes fully aware of their true present bias as soon as $c = \underline{c}$ is realized. If $\beta = \underline{\beta}$, then the agent also becomes fully aware as soon as $c = \underline{c}$, and then switches to no commitment.

Overall, this potential learning channel can explain why agents choose commitment devices that "fail", since this is the only way to update about beliefs. Additionally, if it is not optimal to experiment, then agents never learn but either never go running or always go running.

2.6 Discussion

I have only briefly looked at the potential for learning about self-control issues but it would be interesting to see how this explanation of a limited potential value from learning could help reconcile some of the other findings in the literature. Having non-learning that does not rely on ego concerns, positive value of self-confidence, biased beliefs, or strategic concerns, is appealing and, if plausible, can explain why these other concerns may be allowed to dominate in other contexts. If having precise and unbiased beliefs has little value in many contexts, then the benefits to having distorted beliefs in other contexts ends up dominating overall.

In my model, agents are facing an intra-personal contracting problem. There is no third party, and therefore no possibility of side transfers. Therefore, the agent cannot be compensated for the fine by a lump sum transfer or similar. This setting is plausible for most small-stake but common tasks where agents might be more likely to demand softer commitments, involving image concerns and psychological costs as potential fines, like telling friends that they plan to run, rather than signing formal contracts.

Though Stickk.com is a third party contractor, this is very similar to their setting: on Stickk, a user defines a task and the criteria for its completion, then picks a punishment which can be a monetary amount donated to a charity that the agent would normally not wish to support (e.g. a political party the user does not agree with) or only chooses softer punishment like having friends or a third party referee monitoring progress. In the latter case, only image concerns are at stake and a user is not compensated for the potential punishment. A potential user is facing exactly this tradeoff in my model: choosing punishment that ensures the task will be completed but potentially completing it for high cost realizations, or setting a lower punishment but expecting to fail with positive probability and thereby incurring the costs. There are no side transfers; the website is not offering users a lump sum to sign a contract and then asking them to pay the fine to the website in case of failure.

If the setting described in this paper is empirically prevalent, and this explains limited commitment uptake, it poses complications for quantifying empirical naïveté. It is not enough to observe that agents do not take up commitment to show they are naive, and more importantly, that agents perform better or complete a task more often under externally imposed commitment rather than self-imposed commitment is not proof that agents are partially naive: they might be perfectly sophisticated and choose no commitment or little commitment because commitment is inefficient for high cost realizations.

2.7 Conclusion

In this paper I use an adaptation of the standard present-bias model (Laibson, 1997 and O'Donoghue and Rabin, 1999) with stochastic costs to give new insights into the (lack of) demand for commitment.Overall, I find that if commitment is socially wasteful, agents are unlikely to be better off with commitment than without. This result relies on the assumption of variance in costs, although stochastic costs are not themselves detrimental. On the contrary - if a task is not worthwhile in expectation, variance in costs may make it possible to have positive utility from completing a task. This means that variance in costs increases agents welfare, at least weakly.

I find that commitment is not necessarily implied by sophistication, and that perfect commitment (commitment that commits an agent to behave exactly as a time-consistent agent) is an exception rather than a general result. Therefore, observing agents not behaving as time-consistent agents does not imply that they are naive. In a one-shot setting, it might be optimal to choose a commitment device where one will incur a fee with positive probability for not completing. Additionally, I find that overconfidence is not necessarily bad, but underconfidence always is, which could help explain why, empirically, overconfidence is prevalent. Lastly, since commitment is not attractive to all agents (even if they were perfectly sophisticated), the value of learning about own self-control problems is small, and opportunities to learn may be scarce if agents do not learn about their type without commitment.

Chapter 3

Procrastination and Learning about Self-Control*

3.1 Introduction

Many people procrastinate filing tax returns, setting up a financial portfolio, finding a better credit card contract, or doing a referee report.[1] In some situations, people keep procrastinating the same or similar tasks again and again. Such "naïveté" about own future self-control appears prevalent in many economic environments, and theoretical and empirical literature has investigated its implications. Most of the theoretical literature, however, assumes (either explicitly or implicitly) that people do not learn about their naïveté over time. This opens up a natural question: why do people not learn about their naïveté?

Building upon the literature on naïveté about self-control developed by O'Donoghue and Rabin (1999; 2001), we investigate whether and when a time-inconsistent agent does not learn about her self-control problem over time. In our basic model, a time-inconsistent agent is initially not sure whether her future selves will be time-consistent or time-inconsistent. Note that she initially underestimates the probability that she is actually time-inconsistent. At the beginning of each period, however, she actively chooses whether or not to learn about her own self-control problem by incurring a non-negative

* This chapter is based on joint work with Takeshi Murooka.

[1] Following O'Donoghue and Rabin (2001) and DellaVigna (2009), we say that an agent "procrastinates" if she plans to do a task in some period with some probability but actually does so with lower probability. For evidence of procrastination, see, for example, Ariely and Wertenbroch (2002), DellaVigna and Malmendier (2006), and Skiba and Tobacman (2008).

cost of learning. We highlight the case in which the agent can perfectly learn about her self-control problem and the cost of learning is zero. Crucially, we assume that learning about own self-control takes time (i.e., an agent has to spend some period of time to update her beliefs) and that the agent cannot commit to future actions.

Our key mechanism is that the agent may procrastinate learning about her own self-control. The intuition is two-fold. First, because learning changes own future actions, a time-inconsistent agent may indirectly incur a cost from learning if changing a future action involves both future benefits (e.g., returns from setting up a financial portfolio) and future costs (e.g., effort cost to set up a financial portfolio). We show that this indirect cost occurs only when an agent has intra-personal conflicts between future selves (e.g., hyperbolic discounting), but not when she has intra-personal conflicts only between her current self and future selves (e.g., quasi-hyperbolic discounting). Second, because the time-inconsistent agent initially underestimates the possibility that she will be time-inconsistent, she underestimates the probability that she will procrastinate in the future, and hence she may prefer to delay the learning opportunity. Note that the logic here is akin to O'Donoghue and Rabin (1999, 2001). As a result, the agent may avoid the opportunity to learn about her own self-control problem, even when the (direct) cost of learning is zero. We also derive conditions in which procrastination of learning and non-completion of a task is a *unique* equilibrium outcome. Our result helps explain why people procrastinate even when they face similar tasks repeatedly.

This non-learning result has the following implications. First, if people have time-inconsistent preferences and cannot commit to future actions they may not learn about their naïveté even when the cost of learning is zero. This result complements the study by Ali (2011) which shows that time-inconsistent people will perfectly learn about their self-control if they can take up a flexible commitment device. Indeed, our results bridge the gap between his theoretical result and the empirical prevalence of naïveté. Second, even when people are inherently aware of their own self-control problem, whenever facing a new task they may prefer to forget or neglect their awareness and take an action as an ignorant self. Perhaps surprisingly, this incentive to forget or neglect own self-control problems does not depend on self-esteem (e.g., anticipatory utility, ego utility, private image, or social image) which is often assumed in the literature (e.g., Kőszegi, 2006, Gottlieb, 2014, 2016).

As an important extension, when an agent has multiple initially-uncertain attributes (e.g., she is uncertain about both her own future self-control and own ability for the task), we derive conditions in which the agent's endogenous learning decisions may be misdirected—she chooses to learn what she should not learn from the long-run perspective, and she chooses not to learn what she should learn from the long-run perspective.

Related Literature. There are several strands of literature that investigate why people do not learn over time. Most closely related to our paper is the literature on strategic ignorance: Carrillo and Mariotti (2000) and Bénabou and Tirole (2002) show how an agent may strategically abstain from a learning opportunity to keep motivating their own future selves to work harder under the presence of self-control problems. Relatedly, Bénabou and Tirole (2004) study how an agent may commit to some personal rule to maintain self-reputation. Although our result on non-learning might look similar, the mechanisms are quite different: an agent in these papers chooses not to learn as a means of internal commitment (and to improve own future payoffs), whereas an agent in our paper does so because she underestimates the probability that she will behave time consistently. Indeed, non-learning in our model is harmful from the agent's long-run perspective because it leads to a non-completion of a task.

Among the literature on self-esteem, Kőszegi (2006) and Gottlieb (2014; 2016) investigate models in which an agent may avoid learning because of the presence of ego utility, private image, or anticipatory utility. In our model, such psychological costs are captured as a direct cost of learning. In contrast to the literature, we show that non-learning can occur even when there is no direct (physical or psychological) cost of learning.

The literature on non-Bayesian updating, such as Benjamin et al. (2016), focuses on cases in which an agent does not update her own beliefs according to Bayes' rule. The literature on selective attention, such as Schwartzstein (2014) and Gagnon-Bartsch et al. (2017), analyzes situations in which an agent systematically does not encode a certain type of signal. In contrast to this literature, we focus on the situation in which an agent can perfectly learn about own naïveté and the cost of learning itself can be zero, but the agent actively chooses not to learn about own naïveté.

Finally, our result on multiple initially-uncertain attributes is related to the recent literature on learning under misspecified models Fudenberg et al. (2017); Heidhues et al. (2018); Hestermann and Le Yaouanq (2017). Building upon and extending this literature, our mechanism based on procrastination explains why an agent's learning decisions can be misdirected even when she can choose to perfectly learn about all uncertain attributes.

Structure. The rest of this paper is organized as follows. Section 3.2 introduces our model and discusses its key assumptions. Section 3.3 analyzes an illustrative model of task completion. Section 3.4 investigates the model with task choice and task completion. Section 3.5 examines extensions where an agent initially has multiple uncertain attributes and can choose to learn about either (or both) of them. Section 3.6 briefly discusses other extensions. Section 3.7 concludes. Proofs are provided in the Appendix.

3.2 Model

This section introduces our basic model. Section 3.2.1 sets up the model. Section 3.2.2 discusses our key assumptions.

3.2.1 Setup

A risk-neutral agent can initiate a task in periods $t = 1, 2, \cdots, T$, where $T \geq 2$ is either finite or infinite. Once she initiates a task, from the next period on she can complete the task. Each task is represented by $x = (c, b)$: completing the task in period t gives a cost $c \geq 0$ in period t and brings a benefit $b \geq 0$ in period $t + 1$. There is no (physical or psychological) cost of initiating a task. The agent can initiate (and complete) at most one task throughout the game.

Let u_t denote the agent's period-t instantaneous utility. There are two types of agents: time-consistent and time-inconsistent ones. The type of each agent is persistent throughout the game. For the time-consistent agent, her total utility in period t is $u_t + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} u_\tau$ with $\delta \in (0, 1)$. For the time-inconsistent agent, her total utility in period t is $u_t +$

$\sum_{\tau=t+1}^{\infty} D(\tau - t)u_\tau$, where $D(\tau - t)$ represents her time discounting in $\tau - t$ periods.[2] We assume that $D(t)$ is decreasing and that there exists $B(t) \in (0, 1]$ for $t \geq 1$ such that $D(t) = \delta^t B(t)$. Intuitively, $B(t)$ represents the time-inconsistent part of the agent's preferences. For example, the standard exponential-discounting function is represented by $B(t) = 1$ for $t \geq 1$, present-biased preferences (i.e., quasi-hyperbolic discounting) are represented by $B(t) = \beta$ with $\beta \in (0, 1)$ for $t \geq 1$, and a "modified" hyperbolic discounting function is represented by $B(t) = \frac{1}{1+rt}$ where $r > 0$.[3]

Both types of agents are initially uncertain about their own future self-control. We assume that all agents share the same initial belief and that their initial belief about their own future time consistency, denoted by $\hat{B}(t)$, is $\hat{B}(t) = B(t)$ with probability $1 - q \in (0, 1]$ and $\hat{B}(t) = 1$ with probability q. In line with the literature on limited cognition and beliefs as assets (Bénabou, 2015; Bénabou and Tirole, 2016), we assume that updating her initial belief about her own self-control requires the agent to actively choose to learn (we substantially discuss this assumption in Section 3.2.2). Formally, in each period the agent can acquire a signal about her own self-control by incurring a cost $m \geq 0$, while sticking to her initial beliefs if she has not acquired the signal. To focus on our main mechanism, we assume that learning is perfect and perpetual: the signal is perfectly informative and the agents become completely sophisticated about their own self-control problems for the rest of the game.

The timing of the game is as follows. If the agent has not previously taken any action, in each period she takes one of the following actions: learning about her own self-control problem, choosing a task, or not doing anything. If the agent has initiated a task, in each period the agent takes one of the following actions: learning about own self-control problem, completing the task, or not doing anything. We assume that if the agent is indifferent between taking some action and not doing anything, then she will take the

[2] See, for example, Echenique et al. (2017) for the analysis under general time preferences. Relatedly, Schweighofer-Kodritsch (2017) analyzes Rubinstein's alternative bargaining problem with general time preferences.

[3] Note that a typical hyperbolic discounting function is defined as $D(t) = \frac{1}{1+rt}$ where $r > 0$, instead of $D(t) = \frac{\delta^t}{1+rt}$.

action.[4]

As an equilibrium concept, we extend the perception-perfect equilibria (O'Donoghue and Rabin, 2001) in that (i) in each continuation game an agent chooses her best response given her belief and that (ii) an agent keeps using the initial prior about her own self-control problem if she has not acquired a signal about own self-control.

3.2.2 Discussion of Key Assumptions

Learning is active. A crucial assumption of our model is that the agent has to actively choose to learn: whenever updating her beliefs, she has to "encode" a signal (e.g., she needs to introspect from own past experience).[5] Hence, our model is different from a classical one in which Bayesian updating happens automatically. It is also different from self-confirming equilibria, as the agent in our model updates her beliefs only when she chooses to do so.[6]

Unless the agent chooses to learn, she does not revise her beliefs and keeps using her prior belief. Hence, different from Bénabou and Tirole (2004), we rule out self-signaling over time. In line with O'Donoghue and Rabin (1999; 2001), we also assume away the possibility that the agent becomes aware of her own self-control by inferring from own current preferences (i.e., we assume away inferences like "because my preferences are time-inconsistent today, I must be a time-inconsistent type in the future").[7]

Learning takes time. Another important assumption is that learning about own self-control takes time, i.e., an agent has to spend some period of time to update her beliefs.

[4] This assumption is generically without loss of generality.

[5] In period $t = 1$, we interpret the learning opportunity as the option for the agent to look back to similar past experiences to infer her self-control problem. In period $t \geq 2$, the agent can try to learn about her own self-control problem from own past actions in the game.

[6] Esponda and Pouzo (2016) provide a framework to analyze agents with misspecified models. Note that our model is also different from their framework because in our model the agent can actively choose whether or not to update her initial belief.

[7] Precisely, the crucial assumption throughout our paper is that, as discussed in Footnote 8 of O'Donoghue and Rabin (2001), a time-inconsistent type is "naive about her own naïveté," and hence cannot revise her beliefs unless she acquires information about her self-control. By contrast, all of our results are qualitatively robust (and often become simpler) if a time-consistent type becomes aware of her future self-control by inferring from the fact that she is currently time-consistent.

In our model, this "time cost" is key in generating an intra-personal conflict between future selves (and hence procrastination of learning). While this assumption is restrictive, it captures the notion that we need to spend time to think about ourselves whenever updating such beliefs.

In contrast, the cost of learning itself can be zero. While we highlight the case in which the learning cost is zero (i.e., $m = 0$), we also investigate a case of positive learning cost (i.e., $m > 0$) which represents a physical recollection cost or a psychological cost of deteriorating self-esteem (Kőszegi, 2006; Gottlieb, 2014, 2016).

No commitment. In the model, the agent cannot commit to any particular future behavior, including future learning decisions. In Section 3.6, we discuss the case in which an agent takes up a future commitment device. If the agent can commit to own future actions, non-learning would not occur in our model with zero learning cost. In this respect, our non-learning result complements the study by Ali (2011) which shows that time-inconsistent people will perfectly learn about their self-control if they can take up a flexible commitment device.

3.3 Task Completion

This section analyzes an illustrative example of task completion: we focus on the case in which the agent has already initiated a task $x = (c, b)$. To shed light on our main mechanism in the simplest manner, we further assume in this section that the agent can work on the task only after acquiring a signal about own self-control.[8] We show that even though both types of agents consider the task worthwhile, and even though the task cannot be completed without acquiring information, the time-inconsistent agent will procrastinate (potentially indefinitely) learning her type, even for $m = 0$. In Section 3.4, we analyze a full model without these assumptions.

[8] One interpretation—as we formally investigate in Section 3.4—is that working on a task without knowing own self-control is costly (e.g., the agent cannot optimize her workload) and the agent does not want to engage in any task before learning about own self-control.

3.3.1 Task Completion: Setup

The timing of this example is as follows. If the agent has not acquired a signal $s_t \in \{0, 1\}$ about her self-control problem, in each period she chooses to either acquire it ($s_t = 1$) or not ($s_t = 0$). If the agent has acquired it, in each period the agent chooses whether to complete the task or not.

Let $k \in \{C, I\}$ be the agent's (true) discounting type where C represents a time-consistent type and I represents a time-inconsistent type. Let $U^k_t(x; \tau - t) = -D(\tau - t)c + D(\tau - t + 1)b$ denote a type-k agent's total utility (not taking into account the learning cost m) is evaluated in period t when she completes task x in period τ. If the agent has already encoded the signal (i.e., $s_{\tau_0} = 1$ for some $\tau_0 < t$), the cost of learning $m \geq 0$ is sunk, and hence the agent's utility only depends on when she will complete the task. Before encoding a signal, the agent's subjective expectation of total utility evaluated in period t when she acquires a signal about her type in period $\tau_0 \geq 0$ by $\hat{U}^k_t(s_{\tau_0} = 1)$.

In what follows, we focus on the most interesting case in which each type of agent prefers to complete the task as soon as possible rather than never:

Assumption 3.1 (Task is worthwhile)**.** *(i)* $U^I(x; 0) > 0$*, (ii)* $\min\{U^C(x; 1), U^I_t(x; 1)\} > m$*.*

Because $U^C(x; 0) \geq U^I(x; 0)$, Assumption 3.1 (i) implies that each type of agent prefers to complete the task right now rather than never. Assumption 3.1 (ii) means that each type of agent prefers to acquire information right now and complete the task in the next period rather than never.[9] Note that under Assumption 3.1 both types agree that the task should be completed. Hence, uncertainty and naïveté about own future type is the only reason for procrastination and for not completing the task.

[9] If Assumption 3.1 (i) does not hold, the time-inconsistent agent never completes the task. In this case, if the time-inconsistent agent would not want the task to be completed in any future period or $U^C(x; 0) < 0$, then obviously the time-inconsistent agent would never encode the signal. Otherwise, the time-inconsistent agent may have an incentive to encode a signal, but still in this case procrastinating information acquisition can occur. Regarding Assumption 3.1 (ii), if $U^I_C(x; 1) < m$, then the time-consistent agent never encodes a signal. By contrast, even if $U^I(x; 1) < m$, the time-inconsistent agent may still acquire information when $U^I(x; 0) > 0$.

3.3.2 Task Completion: Analysis in a Finite Period Model

In this subsection, we analyze when procrastination will occur in a task completion setting under a finite time horizon (i.e., $T < \infty$). This corresponds to environments in which there is a deadline for task completion (e.g., filing a tax return).

Since the agent has to learn her type before completing the task, the earliest point in time the task can be completed is $t = 2$, if information was acquired at $t = 1$. Because of the finite time horizon, the equilibria can be identified through backward induction. The agent's decision problem is separated into the optimal timing of information acquisition and the optimal time to complete the task, given the agent's beliefs. In the following, therefore, we analyze when the two types of agents choose to complete the task conditional on having acquired information, and when they acquire information. Because we use backward induction to pin down our equilibrium, we start by characterizing when the different types of agents complete the task, conditional on already knowing their type.

Decisions after learning. By Assumption 3.1, the task is worthwhile for both types of agents. Therefore, due to the properties of a finite horizon decision problem, if the agents have learned their type, they will complete the task at the latest in the last period. Let $\underline{\tau} = \min\{\tau | U^I(x; 0) \geq U^I(x; \tau)\}$ denote the earliest period in which the agent prefers to complete task immediately rather than in $\underline{\tau}$ periods. By Assumption 3.1, both types of agents complete the task in $t = T$ if they have not yet done so. When the task will be completed is summarized in the following lemma:

Lemma 3.1. *Suppose $T < \infty$ and Assumption 3.1 holds. Consider a continuation game in which the agent has already learned her type. Then,*

(i) If the agent is time-consistent, she completes the task immediately.

(ii) If the agent is time-inconsistent, she completes the task in the first period $t \geq \tau$ that satisfies $t = T - n\underline{\tau}$ where $n \in \mathbb{N}_0$.

Note that $U^C(x; k) = \delta^k U^C(x; 0)$ for the time-consistent type. Hence, in any period t, the agent will always complete the task immediately if she is time consistent and the task has not yet been completed.

Lemma 3.1 (ii) implies that a time-inconsistent type completes the task in any $\underline{\tau}$ period. This result is based on O'Donoghue and Rabin (1999; 2001) and its intuition is as follows. Because she is perfectly sophisticated about her future self-control after learning, she knows she will complete the task in the last period and therefore not complete it in the second to last period if $U^I(x;0) < U^I(x;1)$. In the third to last period, she will complete the task if she prefers completing it immediately over completing it in two periods (i.e., $U^I(x;0) \geq U^I(x;2)$); otherwise, she will not. By induction, the agent completes the task in any period $t = T - n\underline{\tau}$ and there is a cycle of length $\underline{\tau}$. Note that this result is not procrastination, as the agent has correct beliefs about when the task will be completed and her actual decisions follow the beliefs.

Information acquisition. Next, we investigate conditions under which each type of agent will acquire information. We focus on deriving the conditions for our equilibrium of interest: an equilibrium in which the time-inconsistent type procrastinates information acquisition and where the time-consistent type acquires information immediately.

We first analyze the actions for the time-consistent type. Note that in this case she is underconfident about own future self-control: she thinks she will be time-inconsistent with probability $1 - q$. If $U^I(x;1) \leq U^I(x;0)$, nevertheless, she will complete the task in the next period with probability one. Otherwise, she takes into account the delay in case her future self would be time-inconsistent and, hence, would complete the task only in period $t = T - n\underline{\tau}$. In the proof, we show that the most stringent condition to acquire information in this case is that the agent expects a delay of $\underline{\tau}$ periods in case her future self would be time inconsistent. The result is summarized in the following lemma:

Lemma 3.2. *Suppose $T < \infty$ and Assumption 3.1 holds. If an agent is time-consistent, she will acquire information in any continuation game if either of the following holds:*

(i) $U^I(x;1) \leq U^I(x;0)$,

(ii) $U^I(x;1) > U^I(x;0)$ *and* $q(1-\delta)U^C(x;1) \geq [1 - \delta q - (1-q)\delta^{\underline{\tau}-1}]m$.

Note that Lemma 3.2 holds when m is close to zero or q is close to one.

Next, we analyze the actions for the time-inconsistent type. First, we illustrate conditions in which the time-inconsistent type postpones acquiring information at $t = T-2$. Because $t = T-1$ is the last period she can acquire information to complete the task, Assumption 3.1 (ii) and Lemma 3.1 ensure that the agent will acquire information at $t = T-1$. Given this, acquiring the information at $t = T-2$ will lead the time-consistent type to complete the task at $t = T-1$, but the time-inconsistent type will postpone the task until $t = T$ if $U^I(x;1) > U^I(x;0)$. Hence, for the time-inconsistent agent in period $t = T-2$, the subjective expected payoff from acquiring the information immediately is

$$\hat{U}^I_{T-2}(s_{T-2} = 1) = -m + qU^I(x;1) + (1-q)U^I(x;2).$$

If the agent postpones acquiring information, she (correctly) anticipates that she will acquire it in the next period and complete the task with certainty in two periods. Hence, the subjective expected payoff from not acquiring information at $t = T-2$ is

$$\hat{U}^I_{T-2}(s_{T-2} = 0) = -mD(1) + U^I(x;2).$$

Therefore, the time-inconsistent agent will prefer not to acquire information at $t = T-2$ if

$$\begin{aligned} & -mD(1) + U^I(x;2) > -m + qU^I(x;1) + (1-q)U^I(x;2) \\ \Leftrightarrow \quad & q(U^I(x;2) - U^I(x;1)) + m(1-D(1)) > 0. \end{aligned} \tag{3.1}$$

Note that she will always prefer to postpone if $U^I(x;2) > U^I(x;1)$. Intuitively, a time-inconsistent agent indirectly incurs a cost from learning because learning itself changes own future actions in this case.

Given Inequality (3.1), we finally show the condition under which the time-inconsistent type *procrastinates* acquiring information at $t = T-2$: In period $t = T-3$, she wrongly believes that she would acquire information at $t = T-2$ with positive probability on the equilibrium path. Suppose that $U^I(x;2) > U^I(x;1) > U^I(x;0)$ and Inequality (ii) holds. Then, upon acquiring information at $t = T-3$, the agent believes that she will complete the task in the next period if and only if her future type will be time-consistent. Hence,

$$\hat{U}^I_{T-3}(s_{T-3} = 1) = -m + qU^I(x; 1) + (1 - q)U^I(x; 3).$$

If she does not acquire information at $t = T - 3$, she (wrongly) anticipates that she will acquire information in the next period with probability q. Hence,

$$\hat{U}^I_{T-3}(s_{T-3} = 0) = q(-mD(1) + U^I(x; 2)) + (1 - q)(-mD(2) + U^I(x; 3)).$$

Note that the beliefs for her future learning decisions are *wrong*: in reality, if she does not learn her own type in period $t = T - 3$, she will not do so in period $t = T - 2$ eiter (because her type is persistent over time). However, she is naive about this exactly because she has not learned about it yet (and hence keeps using her initial prior belief).

Combining these two conditions, the agent will not acquire information at $t = T - 3$, while wrongly believing that she would acquire information at $t = T - 2$ with probability q, if

$$q(U^I(x; 2) - U^I(x; 1)) + m(1 - qD(1)) > m(1 - q)D(2). \tag{3.2}$$

Note that here the logic of procrastination — a time-inconsistent type overestimates own future self-control (in our model, probabilistically) and hence not taking an action now — is akin to the one in O'Donoghue and Rabin (1999, 2001).

The logic and derivations are the same in periods $t < T - 3$, although there are additional conditions depending on $\underline{\tau}$. The result for procrastination until $t = T - 1$ is summarized as follows:

Proposition 3.1. *Suppose $T < \infty$ and Assumption 3.1 holds. If an agent is time-inconsistent, she will procrastinate learning about own future self-control until $t = T - 1$ if $U^I(x; 2) > U^I(x; 1) > U^I(x; 0)$, Inequality* (ii)*, and the following holds:*

$$q(U^I(x; 2) - U^I(x; 1)) + m(1 - qD(1)) > \max_{n,k \in \mathbb{N}} (1 - q)[m \cdot D(n\underline{\tau}) - U^I(x; n\underline{\tau} + k) + U^I(x; k)].$$

Moreover, this equilibrium outcome is unique.

To see the intuition, suppose that $\underline{\tau} = 1$, i.e., a time-consistent type always chooses to acquire information and completes the task in every period. Intuitively, because a time-consistent type is (probabilistically) overconfident about own future self-control, she erroneously thinks that she will acquire information with probability q in the next period, even when she does not do so now. Note that the time-inconsistent type evaluates her anticipated future outcomes with her current (i.e., time-inconsistent) preferences. Hence, as in O'Donoghue and Rabin (1999, 2001), the time-inconsistent agent may procrastinate information acquisition — in any period $t < T - 1$, she thinks that she would acquire information with probability q in the next period, but actually will not do so with probability one.

Perhaps surprisingly, and beyond the original logic of O'Donoghue and Rabin (1999, 2001), our procrastination can occur even when $m = 0$. This is because a time-inconsistent agent indirectly incurs a cost from learning if $U^I(x; 2) > U^I(x; 1)$, unlike for quasi-hyperbolic discounting. Corollary 3.1 highlights the result:

Corollary 3.1. *Suppose $T < \infty$ and $U^I(x; 2) > U^I(x; 1) > U^I(x; 0)$. Then, for any $m \geq 0$, there exists a $\bar{q} < 1$ such that for any $q \in [\bar{q}, 1)$, the time-consistent agent will acquire information immediately while the time-inconsistent agent will procrastinate acquiring information until $t = T - 1$ is a unique equilibrium outcome.*

If the time-inconsistent agent is sufficiently naive (q sufficiently close to 1), she believes that it is of little relevance to her when a time-inconsistent agent would acquire information or complete the task. The only important concern is the behavior of a time-consistent agent, but a time-consistent agent will always acquire information immediately and complete the task immediately upon learning her type. Therefore, the time-inconsistent agent believes that if she acquires information now, she is almost certainly going to complete the task tomorrow (because q is close to one, and hence, she believes she will most likely be time-consistent). If she does not acquire information, she believes that she will most likely acquire it in the next period and then complete the task in two periods. The naive time-consistent agent (erroneously) believes that this is the main trade-off she is facing:

acquiring information today and completing the task tomorrow, or acquiring information tomorrow and completing the task in two periods. If she prefers completing the task in two periods over completing it in one period, she will postpone information acquisition — leading to procrastination.[10]

3.3.3 Analysis: Infinite Horizon Model

In this subsection, we analyze under which conditions a time-inconsistent agent will procrastinate learning information indefinitely, which implies never completing the task on the equilibrium path.

We first specify the agent's behavior in continuation games in which the agent has already learned her type. If the agent learned that she is time-consistent, she chooses to complete the task in any period because Assumption 3.1 (i) implies $U_t^C(x;0) = -c + \delta b > 0$. If the agent learned that she is time-inconsistent, Assumption 3.1 (i) ensures that there exists an equilibrium in continuation games in which the agent completes the task in the next period after learning. When showing the existence of a non-learning equilibrium, we focus on such an equilibrium in continuation games. We then derive the condition in which non-learning for time-inconsistent agents is a unique equilibrium outcome.

Next, we investigate the agent's behavior regarding learning. First, suppose the agent is time-consistent. In this case, Assumption 3.1 (ii) ensures that she chooses to learn about own self-control. Second, suppose the agent is time-inconsistent, note that her beliefs about own future actions depend on the beliefs about own future self-control. On the one hand, the agent (wrongly) believes that she will be time-consistent and hence choose to learn in the next period with probability q. On the other hand, she (correctly) anticipates that she will not choose to learn in any future period with probability $1 - q$ in our candidate equilibrium. To show the existence of the equilibrium, we first select an equilibrium in continuation games in which if the time-inconsistent agent learns about own self-control in period t, then she completes the task in period $t + 1$. Given these

[10] What the time-inconsistent agent fails to realize is that, because she prefers to postpone today she in fact time-inconsistent and will therefore also be time-inconsistent tomorrow. In every period, she is considering the same trade-off, reaching the same conclusion, and taking the same action: postponing until the second to last period.

beliefs, if the agent chooses not to learn, her anticipated expected utility is:

$$\hat{U}_t^I(s_t = 1) = q\left[-D(1)m + U^I(x;2)\right] + (1-q)\cdot 0.$$

If the agent chooses to learn, her anticipated expected utility is:

$$\hat{U}_t^I(s_t = 0) = -m + U^I(x;1).$$

Hence, the agent prefers not to learn if:

$$\begin{aligned} &\hat{U}_t^I(s_t = 0) > \hat{U}_t^I(s_t = 1) \\ \Leftrightarrow\ & q\left[U^I(x;2) - U^I(x;1)\right] + m\left[1 - qD(1)\right] > (1-q)U^I(x;1). \end{aligned} \tag{3.3}$$

Next, we derive the condition under which time-inconsistent agents never choose to learn about their self-control problems—and hence never complete the task—in any equilibrium outcome. Denote the maximum and minimum utility of a time-inconsistent agent from completing a task by $\overline{U}^I(x) = \max_\tau U^I(x;\tau)$ and $\underline{U}^I(x) = \min_\tau U^I(x;\tau)$, respectively. Note that $\overline{U}^I(x) > 0 \geq \underline{U}^I(x)$ by Assumption 3.1 and $\lim_{\tau\to\infty} U^I(x;\tau) = 0$. Then, a lower bound of anticipated expected utility when time-inconsistent agents choose not to learn is:

$$\underline{U}_t^I(s_t = 0) = q\left[-D(1)m + U^I(x;2)\right] + (1-q)\left[-D(1)m + \underline{U}^I(x)\right].$$

The agent knows how a time-consistent agent will behave. Hence, with probability q, she believes she will be time-consistent in the next period, acquire the information, and complete the task in two periods. But with probability $1-q$ she believes she will be time-inconsistent in the next period. In this case, the lowest possible payoff is to acquire information immediately, but only completing the task in the least desired period.

Similarly, an upper bound of anticipated expected utility when the time-inconsistent agents choose to learn is:

$$\overline{U}_t^I(s_t = 1) = -m + qU^I(x;1) + (1-q)\overline{U}^I(x).$$

Again, the action of the time-consistent agent is certain, but, with probability $1-q$, the agent believes she will be time-inconsistent in the next period.

Hence, in any equilibrium outcome, the agent prefers not to learn if:

$$q\left[U^I(x;2) - U^I(x;1)\right] + m\left[1 - D(1)\right] > (1-q)\left[\overline{U}^I(x) - \underline{U}^I(x)\right]. \tag{3.4}$$

The result is summarized as follows.

Proposition 3.2. *Suppose Assumption 3.1 holds.*

(i) If Inequality (3.3) *holds, there exists an equilibrium in which a time-inconsistent agent never learns about own future self-control (and hence never completes the task).*

(ii) If Inequality (3.4) *holds, a time-inconsistent agent never learns about own future self-control in any equilibrium outcome (and hence never completes the task) .*

The intuition of the result is two-fold. First, because the agent is time-inconsistent, she may prefer to complete a task later rather than sooner. Specifically, if the agent prefers to complete the task in two periods rather than in the next period (i.e., if $U^I(x;2) > U^I(x;1)$), then the agent has an incentive to postpone acquiring the signal in order to delay task completion. Second, because the agent is (probabilistically) overconfident about own future self-control, she underestimates the likelihood that she will not acquire the signal in the future. Specifically, if the degree of naïveté (i.e., q) is sufficiently large, the agent (wrongly) believes that she will most likely acquire the signal in the next period and, hence, she prefers not acquiring it now to delay task completion. The time-inconsistent agent thereby fails to infer from her own actions that she cannot be time-consistent in the next period and, hence, she repeatedly procrastinates.

It is worth mentioning that procrastination of learning about self-control can occur even when the time-inconsistent agent prefers to do the task in the next period rather than never (i.e., $U^I(x;1) > m$). Intuitively, because the agent is overconfident about own future self-control, she underestimates the possibility that she will procrastinate learning in the future. As a result, she may prefer to postpone a valuable learning opportunity again and again. Furthermore, if the agent is sufficiently naive and has preferences for a delay, procrastination of learning self-control is a unique equilibrium outcome:

Corollary 3.2. *Suppose Assumption 3.1 holds and $U^I(x;2) > U^I(x;1)$. Then, there exists $\bar{q} \in (0,1)$ such that for any $q > \bar{q}$ a time-inconsistent agent never learns about own future self-control (and hence never completes the task) in any equilibrium outcome.*

Perhaps surprisingly, even when $m = 0$, not learning about self-control can occur in equilibrium. Further, even when $m = 0$, non-learning self-control can be a unique equilibrium outcome.

As an illustrative example, we discuss the case in which the agent's discount function is (modified) hyperbolic:

Example 3.1. *Consider the case in which $D(t) = \frac{1}{1+rt}\delta^t$, $r = \frac{1}{2}$, and $\delta \simeq 1$.*

If $m = 0$, Assumption 3.1 holds if and only if $\frac{b}{c} > \frac{3}{2}$. In this case, Inequality (3.3) becomes:

$$q > \frac{-20c + 15b}{-15c + 12b}.$$

When $\frac{b}{c} = \frac{23}{15}$, the condition under which procrastination can occur is $q > \frac{15}{17}$. Also, when $\frac{b}{c} = \frac{23}{15}$ and $q > \frac{15}{17}$, $\overline{U}^I(x) = U_t^I(x;2)$ and $\underline{U}^I(x) = 0$. Hence, by Inequality (3.4), a time-inconsistent agent never learns about own future self-control (and hence never completes the task) in any equilibrium outcome under these parameters.

If $\frac{m}{c} = \frac{1}{3}$, Assumption 3.1 holds if and only if $\frac{b}{c} > 2$. In this case, Inequality (3.3) becomes:

$$q > \frac{-90c + 45b}{-65c + 36b}.$$

When $\frac{b}{c} = \frac{91}{45}$, the condition under which procrastination can occur is $q > \frac{5}{39}$.

3.4 Task Choice

So far, we have assumed that the agent already faces a specific task to complete. This section analyzes a situation where the agent has multiple tasks to choose from and where the optimal choice depends on the agent's type. For example, such a task could be to choose between different financial contracts. As an illustrative example, we will focus on a case where the agent is facing two possible tasks to choose from.

3.4.1 Analysis

Suppose that the agent can choose one of the following two tasks: $x = (c, b)$ and $x' = (0, b')$ with $-c + \delta b > \delta b' > 0 > -c + D(1)b$. As described in Section 3.2.1, we assume that the task cannot be completed in the same period it is chosen. Unlike the illustrative model analyzed in Section 3.3, however, in this full model we allow the agent to choose a task without knowing her self-control problem.

We derive the conditions under which time-inconsistent agents never choose to learn about own self-control problems nor choose a task in equilibrium.

First, we characterize the agent's task-completion behavior. It is straightforward to show that (i) if the agent takes up task x, she will choose to complete the task in every period if she is time-consistent and will never complete the task if she is time-inconsistent, (ii) if the agent takes up task x', the agent will choose to complete the task in every period irrespective of her type.

Second, suppose the agent has learned about her self-control, but has not yet chosen a task. In this case, it is straightforward to show that the agent would choose x in any period if she is time-consistent and would choose x' in any period if she is time-inconsistent.

Third, suppose that the agent has neither learned about own self-control nor chosen a task. If she is time-consistent, she strictly prefers learning about own self-control rather than not doing anything if m is small (i.e., $m < \min\{\delta^3 b', U^I(x; 2)\}$).[11] The perceived

[11] Intuitively, cquiring information immediately implies incurring cost m with certainty, but only completing the task x in the future with perceived probability q, while doing nothing implies the information cost will only be incurred if the agent is in fact time-consistent.

expected utility of a time-consistent type for each action is as follows:

- Taking x without learning self-control: $qU^C(x;1) + (1-q)\cdot 0$.
- Taking x' without learning self-control: $U^C(x';1)$.
- Learning self-control: $-m + qU^C(x;2) + (1-q)U^C(x';2)$.
- Not doing anything: $q(-m\delta + U^C(x;3)) + (1-q)\cdot 0$.

Hence, if the agent is time-consistent, she will choose to learn if and only if:

$$-m + qU^C(x;2) + (1-q)U^C(x';2) > \max\{qU^C(x;1), U^C(x';1)\}$$

or

$$\frac{(1-q)\delta}{1-\delta}U^C(x';0) - \frac{m}{\delta(1-\delta)} > qU^C(x;0) > \frac{1-(1-q)\delta}{\delta}U^C(x';0) + \frac{m}{\delta^2} \tag{3.5}$$

Note that this always holds for $m = 0$ if δ is close to one.

Consider the case in which Inequality (3.5) holds. Suppose that the agent has neither learned about own self-control nor chosen a task and that the agent is time-inconsistent. If she chooses to learn and she finds out that she will have time-consistent preferences in the future, she expects that she will take up task x in the next period and then complete it in two periods. Her perceived expected utility for each action is as follows:

- Taking x without learning self-control: $qU^I(x;1) + (1-q)\cdot 0$.
- Taking x' without learning self-control: $U^I(x';1)$.
- Learning self-control: $-m + qU^I(x;2) + (1-q)U^I(x';2)$.
- Not doing anything: $-mqD(1) + qU^I(x;3) + (1-q)\cdot 0$.

Hence, if the agent is time inconsistent, she will choose not to do anything if

$$U^I(x;3) - mD(1) > \max\left\{U^I(x;1),\ \frac{1}{q}U^I(x';1),\ U^I(x;2) + \frac{1-q}{q}U^I(x';2) - \frac{m}{q}\right\} \tag{3.6}$$

and Inequality (3.5) hold.

The intuition here is threefold. First, if taking up a task without learning own self-control is costly, the agent thinks that she will acquire the signal about own self-control if she will be time-consistent (this happens whenever the agent has a sufficiently large δ). Second, because the agent is time-inconsistent, she may prefer to complete a task later rather than sooner. Specifically, if $U^I(x;3) > \max\{U^I(x;1), U^I(x;2)\}$, the agent has an incentive to postpone acquiring the signal to delay task completion. Third, because the agent is (probabilistically) overconfident about own future self-control, she underestimates the likelihood that she will not acquire the signal in the future. The next proposition summarizes our main result:

Proposition 3.3. *Consider the task choice model. If $m < \min\{\delta^3 b', U^I(x;2)\}$, Inequality (3.5), and Inequality* (3.6) *hold, there exists an equilibrium in which a time-inconsistent agent never learns about own future self-control nor choose any task.*

Proposition 3.3 highlights a perverse welfare effect of non-learning in our model: If the inconsistent type would know that she is inconsistent, then she would choose task x' which gives her strictly positive utility. But because she believes that she is time-consistent with a high probability, she believes that she will choose task x in the future with a high probability, and therefore, she does nothing indefinitely.

As a real-world example, suppose that an agent has to choose whether to pursue a university degree or do an apprenticeship. If she is time-consistent, she would go to university this year. If she is time-inconsistent, the high upfront cost of the university degree is too much for her, so she would choose the apprenticeship this year. But if she is time-inconsistent and does not know her future type, she may believe that her future selves would be time-consistent with a high probability and hence go to university next year. This is why she does not start the apprenticeship this year. In fact, she will never complete any further education.

In the following, we will use our preferred functional form of discounting $D(t) = \frac{1}{1+rt}$ to show under which degree of naïveté q and under which parameter restrictions we can observe an equilibrium where the time-inconsistent agent never chooses or completes a task.

First, we look at the case where the cost of acquiring information is zero.

Example 3.2. *Consider the case in which* $D(t) = \frac{1}{1+rt}$, $r = \frac{1}{2}$, $m = 0$, *and* $\delta \simeq 1$. *Then, the assumption* $-c + \delta b > 0 > -c + D(1)b$ *holds if and only if* $\frac{b}{c} \in (1, \frac{3}{2})$. *In this case, Inequality 3.5 always holds. Inequality ((3.6)) is equivalent to the following condition:*

$$q > \max\left\{\frac{25b'}{c + 20b'}, \frac{15b'}{4c}\right\}. \tag{3.7}$$

Hence, for $\frac{b'}{c} < \frac{1}{5}$ *there exists a* $\bar{q} < 1$ *such that for all* $q > \bar{q}$ *a time-inconsistent agent will never learn about own self-control and never initiate a task.*

We already know that if $m = 0$ and $\delta = 1$, then the time-consistent agent will always choose to learn if she has taken no other action, and will always complete her chosen task immediately if she has already initiated a task. Therefore, we only need to ensure that inequality (3.6) is satisfied. What this example shows is that for (modified) hyperbolic discounting, a sufficiently naive agent (q close to one) will procrastinate information acquisition and task choice indefinitely if b' is small. The intuition is as follows: if q is large, the agent (wrongly) believes that it is likely that x is the optimal task. Therefore, taking up task x' becomes unattractive. But since the agent is time-inconsistent she might wish to postpone the task, and taking up the task x immediately will lead to completion in the next period with perceived probability q. Therefore, not doing anything is the only way to commit to postpone the task completion, according to her beliefs.

In the example, if b' approaches zero, the degree of naïveté q ensuring procrastination goes to zero. This means that when the value of knowing that she is time-inconsistent goes down, she becomes more likely to procrastinate and stick to a long-run suboptimal decision. Similarly, if c increases while $\frac{b}{c}$ is fixed, the time-inconsistent agent is more likely to procrastinate. As in the task completion setting, the time-inconsistent agent is expecting to (most likely) be time-consistent in the future, but she is evaluating the task x according to her own preferences — the preference of a time-inconsistent agent. Despite knowing that a time-inconsistent agent would never complete the task x, she is still attempting to commit herself to completing the task in three periods.

In the next example, we look at what happens with positive information cost m.

Example 3.3. *Consider the case in which* $D(t) = \frac{1}{1+rt}$*,* $r = \frac{1}{2}$*,* $\delta \simeq 1$*,* $m = \frac{c}{50}$*, and* $\frac{b'}{c} = \frac{1}{10}$*. Then, the assumption* $-c + \delta b > 0 > -c + D(1)b$ *holds if and only if* $\frac{b}{c} \in (1, \frac{3}{2})$*. In this case, Inequality* (3.5) *and* (3.6) *become:*

$$\frac{4}{5} > q > \max\left\{\frac{1}{50\frac{b}{c} - 55}, \frac{15}{44}\right\}. \tag{3.8}$$

Hence, for $\frac{b}{c} \in (\frac{9}{8}, \frac{3}{2})$*, there will is a range for* q *such that the time-consistent agent acquires information, and the time-inconsistent agent never learns about own self-control and never completes any of the tasks.*

With a strictly positive information cost m, the time-consistent agent may choose to pick a task without having acquired information. This poses restrictions on the degree of naïveté q. Intuitively, if the agent believes she is particular type with sufficient certainty (perhaps wrongly), the information has low perceived value and the agent will pick the task she believes most likely to match her type. If the net value of the task x for the time-consistent agent is sufficiently low, the condition (3.5) becomes binding. This means that to ensure the time-consistent agent prefers acquiring information is a stronger requirement on q than to have the time-inconsistent agent procrastinate.

If q is sufficiently high such that the time-consistent agent picks x without acquiring information, the time-inconsistent agent no longer believes she can postpone the task for as long as before. In this case, she might pick x' to ensure the task x is not completed too soon, resulting in a perceived negative payoff.

3.5 Procrastination and Misdirected Learning

This section investigates an extension to our basic model by including learning about an additional payoff-relevant attribute: ability. Suppose that there are two initially-uncertain attributes: self-control and ability. The agent is initially overconfident and partially naive. The agent can acquire perfectly informative signals about either of the two factors (or both; but not both in one period).

As in our basic model, the agent is partially naive, but since we now have two factors the agent is uninformed about, we need further notation. Let q_d be the probability with which the agent believes her discounting is exponential, i.e. the probability that she is time-consistent, and let $(1 - q_d)$ be the probability the agent attaches to having discount function $D(t)$, displaying time-inconsistent preferences. Additionally, the agent now has beliefs about her ability. For simplicity we assume that there are only two ability types: high ability, a_H, and low ability, a_L. The agent has beliefs $\hat{a}$ about her ability a: she believes with probability q_a that she is of high ability, a_H, and with probability $(1 - q_a)$ that she is of low ability, a_L. Let $m_a \in \mathbb{R}$ denote the cost of learning own ability. Where m_a can be negative.

3.5.1 Illustration

For illustrative purposes, we first examine a version of the model in which the agent has an opportunity to learn about own ability only in $t = 0$, and then plays the game described in Section 3.3. That is, the agent faces a task x with payoff (c, b_i) where $b_i = a_i b$ depends on ability. The agent, therefore, believes her expected ability is $\hat{a} = qa_H + (1 - q)a_L$, and the expected delayed benefit of the task is $\hat{a}b = \hat{b}$. With slight abuse of notation, let $U_t^k(b_i; \tau - t) = -D(\tau - t)c + D(\tau - t + 1)b_i$ denote a type-k agent's total utility (not taking into account the learning cost m) evaluated in period t when she completes task x in period τ.

First, consider the case in which $\max_\tau U_t^I(\hat{b}; \tau) = U_t^I(\hat{b}; 1)$, $U_t^C(b_L; 1) > 0$, and $U_t^I(b_L; 2) > 0 > U_t^I(b_L; 1)$. In this case, a time-inconsistent type never works on the task once she learns that her ability is a_L. However, from her $t = 0$ perspective, she wants to complete the task. Also, because $\max_\tau U_t^I(\hat{b}; \tau) = U_t^I(\hat{b}; 1)$ (i.e., b_H is sufficiently large), without learning own ability, even a time-inconsistent type would immediately learn own self-control and then complete the task. That is, the agent would strictly prefer to avoid learning own ability to motivate her future self to complete the task, as analyzed by Carrillo and Mariotti (2000) and Bénabou and Tirole (2002).

Second, consider the case in which $U_t^I(b_H; 1) > 0 > U_t^C(b_L; 1)$. This implies that $U_t^C(b_H; 1) > 0 > U_t^I(b_L; 1)$. In this case, each type of agent never works on the task once she learns that her ability is a_L. Since information regarding ability is valuable, a time-consistent type chooses to learn at $t = 0$. As $q \to 1$, a time-inconsistent type also (wrongly) believes that, almost surely, she will behave as if she would be time-consistent. Precisely, if Inequality (3.3) holds, then both time-consistent and time-inconsistent types strictly prefer to acquire information about own ability at $t = 0$.

The following proposition summarizes the results:

Proposition 3.4. *Suppose $T = \infty$ and the agent has an opportunity to learn own ability only in $t = 0$.*

(i) Assume $\max_\tau U_t^I(\hat{b};\tau) = U_t^I(\hat{b};1)$, $U_t^C(b_L;1) > 0$, $U_t^I(b_L;2) > 0 > U_t^I(b_L;1)$, and $m_a = 0$. If q is sufficiently close to 0, there exists an equilibrium in which a time-inconsistent agent strictly prefers to not learn own ability and then completes the task at $t = 2$.

(ii) Assume $U_t^I(b_H;1) > 0 > U_t^C(b_L;1)$ and $m_a = 0$. If Inequality (3.3) holds, there exists an equilibrium in which a time-inconsistent agent strictly prefers to learn about own ability, but then never learns about own future self-control nor completes the task.

Although our result may look close to the literature on strategic ignorance (Carrillo and Mariotti, 2000; Bénabou and Tirole, 2002), the mechanisms are quite different. In the literature, an agent chooses not to learn about own ability as a means of internal commitment to improve own future payoffs. In this sense, the result in Proposition 3.4 (i) is in line with the literature.

By contrast, in Proposition 3.4 (ii) the agent chooses not to learn about own future self-control because she is naive about own future self-control and hence procrastinates learning. Indeed, non-learning about own self-control in our model is harmful from the agent's long-run perspective because it leads to a non-completion of a task. In this case, the agent also acquires information about own ability even though it never benefits her (indeed, acquiring it strictly lowers the agent's long-run utility). In this sense, Proposition 3.4 (ii) highlights how the agent's endogenous learning decisions may be misdirected—she chooses to learn what she should not learn from the long-run perspective, and she chooses not to learn what she should.

3.6 Extensions and Discussion

Another way of modeling (partial) naïveté is having preferences being stochastic in each period, i.e., the true probability of being time-inconsistent in each period is q and its realization is i.i.d. across time, but the agent anticipates that it is $\hat{q}$. In the limit case where $q = 0$, we are in the same world as our model, but with a different form of (partial) naïveté. This means that the optimal decision given full information has not changed, only the perceived optimal action from the perspective of partially naive agents.

We have briefly discussed the implications of allowing for commitment in our model. In the task-completion model, both agents will immediately commit to completing the task in their preferred period. For the time-consistent type, she commits to complete the task as soon as possible. For the time-inconsistent agent type, she commits to complete the task in the future period that maximized the payoff for the agent from the perspective of the current period. From Assumption 3.1, we know the time-inconsistent agent will have a positive utility from completing the task when she completes it. As is described, the non-commitment assumption is crucial in our results.

Because the non-learning in our model can occur even for $m = 0$, they may also occur for some small $m < 0$. It implies that agents may be willing to pay some cost to stay ignorant, or will forgo small positive payoffs from learning her type today.

3.7 Concluding Remarks

This paper provides a new mechanism for why people do not learn about their naïveté over time. We find that individuals may procrastinate a free learning opportunity indefinitely, even though having the information would make the agent better off. We also find that the more biased a time-inconsistent agent's beliefs are, the more likely she is to stay naive. Our results do not rely on agents having image concerns or on using overconfidence as commitment. Indeed, naïveté in our model has no benefits to the agent and hence is purely harmful.

One potential future direction of research is applying it to situations which involve strategic interactions between players. For example, when firms offer contracts which specify a base contract and a set of options, they often have an incentive to make consumers procrastinate canceling or switching the option when a fraction of consumers are naive (DellaVigna and Malmendier, 2004; Heidhues and Kőszegi, 2010; Murooka and Schwarz, 2018). Our results imply that firms may have an incentive to make consumers endogenously procrastinate learning about own naïveté by setting an appropriate contract and pricing structures. How the procrastination of learning about own naïveté can be interacted with strategic concerns of other parties is left for future research.

Appendix A

Packaging of Sin Goods - Commitment or Exploitation?

Relaxing the assumption of observable degree of sophistication (monopolistic case)

In the main text, we assumed that the monopolist can distinguish between the naive and the sophisticated consumer. We now discuss what happens if we relax this assumption. Suppose first that the monopolist caters to the naive self 1 by offering (x_1^*, t_1^*). Then $v(x_1^*) - c\,x_1^* - k_1 \geq u(x_0^*) - c\,x_0^*$.

The sophisticated consumer pays transfer $u(x_0^*) - u(x_1^*) + v(x_1^*)$ for x_0^*. Similarly, he would pay $u(x_0^S) - u(x_1^*) + v(x_1^*)$ for x_0^S. The naive consumer has no willingness to pay for commitment. Thus, the naive self 0 would not choose x_0^* or x_0^S as the associated transfer (which extracts the sophisticated consumer's willingness to pay for commitment) is too high. But the sophisticated consumer would prefer buying (x_1^*, t_1^*) over buying x_0^* or x_0^S. The transfers for these quantities are constructed such that the sophisticated consumer is indifferent between buying at time 0 and buying x_1^* at time 1. Therefore, the sophisticated consumer is paying to avoid the higher shopping cost k_1. If the sophisticated consumer can not be identified, and hence can buy x_1^* at time 0, the consumer can avoid the shopping cost k_1 and hence keep rents from the transaction. In order to avoid the sophisticated consumer imitating the naive, the monopolist will either have to subtract k_1 from the time 0 transfer to again make the sophisticated consumers indifferent, offer the same bundle to all consumers, or exclude the naive consumers.

Suppose next the monopolist caters to the naive self 0. Then $u(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$. Note that then $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$ also holds. Hence, the sophisticated consumer also receives x_0^*. But since the monopolist prefers to offer x_0^* at a higher transfer to the sophisticated consumer than to the naive consumer the monopolist either offers only (x_0^*, t_0^S) and serves only the sophisticated consumers or offers (x_0^*, t_0^*) and serves both types or offers both $(x_0^*, t_0^S - k_1)$ and (x_1^*, t_1^*) again offering both types. What is optimal depends on the share of naive and sophisticated consumers, and the shopping cost k_1. If, for example, the share of sophisticated consumers is large, he prefers to offer x_0^* at the higher transfer and not serve the naive consumers. If the share is small he prefers serving everyone the same quantity-transfer pair, but there may exist an intermediate share where it is optimal to serve sophisticated consumers $(x_0^*, t_0^S - k_1)$, and the naive (x_1^*, t_1^*). In this case the loss in profits from subtracting the shopping cost k_1 to make the sophisticated consumers indifferent, is smaller than the gain in profits from serving naive consumers at time 1.

Suppose $u(x_0^*) - c\,x_0^* \leq v(x_1^*) - c\,x_1^* - k_1$, then the monopolist will either exclude naive consumers and serve sophisticated consumers (x_0^*, t_0^S) or (x_0^S, t_0^S), or serve sophisticated consumers $(x_0^*, t_0^S - k_1)$ or $(x_0^S, t_0^S - k_1)$ and serve the naive consumers (x_1^*, t_1^*).

Proposition A.1. *Suppose the monopolist cannot distinguish naive and sophisticated consumers.*

1. *Suppose $u(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$. Then*
   - *(a) For large shares of naive consumers the monopolist will offer (x_0^*, t_0^*) to all consumers.*
   - *(b) For small shares of naive consumers (x_0^*, t_0^S) to sophisticated consumers, excluding naive consumers.*
   - *(c) For intermediate share of naive consumers potentially offer $(x_0^*, t_0^S - k_1)$ to sophisticated consumers and (x_1^*, t_0^*) to naive consumers.*
2. *Suppose $v(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1 \geq u(x_0^*) - c\,x_0^*$. Then*
   - *(a) For large shares of naive consumers the monopolist will offer $(x_0^*, t_0^S - k_1)$ to sophisticated consumers and (x_1^*, t_1^*) to naive consumers.*

(b) *For small shares of naive consumers the monopolist will offer* (x_0^*, t_0^S) *to sophisticated consumers and exclude naive consumers.*

3. *Suppose* $v(x_1^*) - c\,x_1^* - k_1 \geq v(x_0^*) - c\,x_0^*$*. Then*

   (a) *For large shares of naive consumers the monopolist will offer* $(x_0^S, t_0^S - k_1)$ *to sophisticated consumers and* (x_1^*, t_1^*) *to naive consumers.*

   (b) *For small shares of naive consumers the monopolist will offer* (x_0^S, t_0^S) *to sophisticated consumers and exclude naive consumers.*

Proof. Assume λ is the share of naive consumers, and $u(x_0^*) - c\,x_0^* \geq v(x_1^*) - c\,x_1^* - k_1$. Then the firms options are: serve everyone at time 0, $(x_0, t_0) = (x_0^*, u(x_0^*))$, only serve sophisticated, $(x_0, t_0) = (x_0^*, u(x_0^*) + v(x_1^*) - u(x_1^*))$, or serve sophisticated at time 0 and naive at time 1, $(x_0, t_0) = (x_0^*, u(x_0^*) + v(x_1^*) - u(x_1^*) - k_1)$, $(x_1, t_1) = (x_1^*, v(x_1^*))$. This gives the following profits for the firm respectively:

Pooling:	$u(x_0^*) - c\,x_0^*$
Sophisticated:	$(1-\lambda)(u(x_0^*) + v(x_1^*) - u(x_1^*) - c\,x_0^*)$
Split:	$\lambda(v(x_1^*) - c\,x_1^*) + (1-\lambda)(u(x_0^*) + v(x_1^*) - u(x_1^*) - k_1 - c\,x_0^*)$

Comparing these profits pairwise yields the following cutoffs:
Split equilibrium over pooling:

$$\lambda \leq \frac{v(x_1^*) - u(x_1^*) - k_1}{u(x_0^*) - c\,x_0^* - (u(x_1^*) - c\,x_1^*)}.$$

Sophisticated over pooling:

$$\lambda \leq \frac{v(x_1^*) - u(x_1^*)}{u(x_0^*) - c\,x_0^* + v(x_1^*) - u(x_1^*)}.$$

Split equilibrium over sophisticated:

$$\lambda \geq \frac{k_1}{v(x_1^*) - c\,x_1^*}.$$

From this we have that there exists a $\lambda \in (0,1)$ such that the monopolist is indifferent between only serving sophisticated and serving the two types at separate times. We also have that $\lambda \in (0,1)$ such that the monopolist is indifferent between serving types at different times and serving everyone at time 0. But it might always be preferred to serve everyone at time 0 over splitting the consumers. If:

$$\frac{v(x_1^*) - u(x_1^*) - k_1}{u(x_0^*) - c\,x_0^* - (u(x_1^*) - c\,x_1^*)} \in (0,1).$$

The monopolist will only serve sophisticated if:

$$\lambda \leq \frac{k_1}{v(x_1^*) - c\,x_1^*}.$$

The monopolist will serve types at different times if:

$$\frac{k_1}{v(x_1^*) - c\,x_1^*} \leq \lambda \leq \frac{v(x_1^*) - u(x_1^*) - k_1}{u(x_0^*) - c\,x_0^* - (u(x_1^*) - c\,x_1^*)}.$$

And serve everyone at time 0 if:

$$\lambda \geq \frac{v(x_1^*) - u(x_1^*) - k_1}{u(x_0^*) - c\,x_0^* - (u(x_1^*) - c\,x_1^*)}.$$

If:

$$\frac{v(x_1^*) - u(x_1^*) - k_1}{u(x_0^*) - c\,x_0^* - (u(x_1^*) - c\,x_1^*)} \notin (0,1)$$

the monopolist will serve sophisticated if:

$$\lambda \leq \frac{v(x_1^*) - u(x_1^*)}{u(x_0^*) - c\,x_0^* + v(x_1^*) - u(x_1^*)}.$$

And serve everyone at time 0 if:

$$\lambda \geq \frac{v(x_1^*) - u(x_1^*)}{u(x_0^*) - c\,x_0^* + v(x_1^*) - u(x_1^*)}.$$

Assume $u(x_0^*) - c\,x_0^* \leq v(x_1^*) - c\,x_1^* - k_1$. Then the pooling equilibrium becomes $(x_1^*, v(x_1^*) - k_1)$, the split equilibrium only changes in that sophisticated consumers now demand $x_0^S \geq x_0^*$. Comparing the profits of pooling and split equilibrium shows that split is always preferred over pooling. Hence there will only be one cutoff, where the monopolist is indifferent between only serving sophisticated consumers, and serving sophisticated at time 0 and naive at time 1. The cutoff is unchanged from before, and hence for:

$$\lambda \geq \frac{k_1}{v(x_1^*) - c\,x_1^*}$$

the monopolist will only serve sophisticated, and otherwise serve consumers at different times. □

Relaxing the assumption of on-the-spot-production (monopolistic case)

The solution for the naive consumer in the monopolistic market does not rely on the assumption that on-the-spot-production is feasible. So consider the sophisticated consumer. If the monopolist cannot produce on the spot, then, when facing a sophisticated consumer, he maximizes his profits subject to the participation constraint of self 0 and/or self 1. The solution coincides with the solution for the naive consumer. So the monopolist does not offer partial commitment to self 0 any longer.

Proof of the contract for the naive self 1 in a competitive market

Suppose first that PC_1 binds. Rewrite it to get $t_1 = v(x_0 + x_1) - k_1 - \max\{v(x_0), \bar{v}^*\}$ and plug t_1 into the objective function. Then we can write the problem of the firm as $\max_{x_1} v(x_0 + x_1) - c\,x_1 - k_1 - \max\{v(x_0), \bar{v}^*\}$. The first order condition is $v'(x_0 + x_1) = c$. Hence, by the definition of x_1^* it follows that $x_0 + x_1 = x_1^*$, i.e., $x_1 = x_1^* - x_0$. Thus, $v^* = v(x_1^*) - k_1 - c\,(x_1^* - x_0)$ and $t_1 = c\,(x_1^* - x_0)$. Finally, note that $\max\{v(x_0), \bar{v}^*\} = v^*$ whenever $v(x_1^*) - c\,x_1^* - k_1 \geq v(x_0) - c\,x_0$.

Package sizes

In the competitive equilibrium, there are either one or three package sizes as described in the main text. For the equilibrium with three quantities (x_0^*, x_0^S and $x_1^* - x_0^*$), we ask how these relate to one another in size. We can rewrite the condition of optimal choice of x_0^S as:

$$k_1 = v(x_1^*) - v(x_0^S) - c(x_1^* - x_0^S)$$

For $k_1 > 0$ we have that $x_0^S \in (x_0^*, x_1^*)$. From this we can see that for k_1 small enough we can get x_0^S arbitrarily close to x_1^*, though $x_0^S < x_1^*$. Since we have assumed three quantities we have $x_0^S > x_0^*$, though for k_1 high enough we can get arbitrarily close to x_0^*.

The question is then how $x_1^* - x_0^*$ relate to the other two quantities. If $x_1^* < 2\,x_0^*$ implying the consumer will want to consume less than twice the amount when present biased compared to when not, the top-up quantity $x_1^* - x_0^*$ is smaller than x_0^* and hence the top-up quantity is the smallest in the market.

When does this arise? Suppose, using the present-bias formulation outlined in footnote 5, that $u(x) = b(x) - \kappa(x)$ and $v(x) = b(x) - \beta\,\kappa(x)$. Intuitively, the larger β, i.e., the less severe the self-control problem, the closer x_1^* is to x_0^* (for $\beta = 1$, $x_1^* = x_0^*$) and thus, the more likely $x_1^* < 2\,x_0^*$ holds. Assuming, e.g., a linear delayed cost function ($k(x) = \kappa\,x$) and a quadratic or logarithmic immediate benefit function ($b(x) = \ln x$ or $b(x) = \sqrt{x}$) shows that $x_1^* < 2\,x_0^*$ arises for all $\beta > \frac{1}{2}$ (and possibly for all $\beta \in (0,1)$, depending on κ and c). For a linear current benefit function ($b(x) = b\,x$) and a quadratic delayed cost function ($\kappa(x) = \frac{x^2}{2}$), $x_1^* < 2\,x_0^*$ arises for $\beta > \frac{1}{2}$. Using a real effort task, Augenblick, Niederle and Sprenger 2013 estimate a β around 0.9. That is, $\beta > \frac{1}{2}$ seems the empirical plausible size of the present bias.

If on the other hand $x_1^* > 2\,x_0^*$ the quantity x_0^* is smaller, and since we can find a k_1 so that x_0^S is arbitrarily close to x_0^*, it is even possible that the top up quantity is the biggest in the market. This will happen if the present bias is severe, so that $x_1^* >> x_0^*$ and k_1 is large such that $x_0^S - x_0^*$ is small.

Proof Proposition 1.2

The proof for the sophisticated consumer is in the text. So consider the naive consumer. Suppose it is optimal to sell some quantity $x_0 > 0$ to self 0 and $x_1 > 0$ to self 1. Then it follows from the participation constraint of self 0 that $t_0 = u(x_0)$. And from the participation constraint of self 1 it follows that $t_1 = v(x_0 + x_1) - k_1 - v(x_0)$. Thus, the monopolist maximizes $v(x_0 + x_1) - v(x_0) + u(x_0) - c\,(x_0 + x_1)$ over x_0 and x_1. The first order conditions are:

$$x_0 : \qquad v'(x_0 + x_1) - v'(x_0) + u'(x_0) \leq c \quad \text{with equality if } x_0 > 0,$$
$$x_1 : \qquad v'(x_0 + x_1) \leq c \quad \text{with equality if } x_1 > 0.$$

The two first order conditions cannot hold with equality at the same time. Hence, either $x_0 > 0$ and $x_1 = 0$, or $x_0 = 0$ and $x_1 > 0$. In the former case the optimal transfer and quantity are determined by:

$$u'(x_0^*) = c \quad \text{and} \quad t_0^* = u(x_0^*).$$

In the latter case, the optimal transfer and quantity are determined by:

$$v'(x_1^*) = c \quad \text{and} \quad t_1^* = v(x_1^*) - k_1$$

Whether it is optimal to cater to the interest of self 0 or self 1 depends on whether $v(x_1^*) - c\,x_1^* - k_1$ is smaller or larger than $u(x_0^*) - c\,x_0^*$.

Partial naïveté

Let $\hat{v}(x) = v_\theta(x) = (1-\theta)\,u(x) + \theta\,v(x)$. Where $\theta \in (0,1)$ is a parameter capturing the severity of naivety.[1] Then if:

$$\nu(x_0^*) - cx_0^* > \nu(x_1^*) - c\,x_1^* - k_1,$$

even perfectly sophisticated consumers will chose x_0^* since this is sufficient to prevent self 1 from shopping. In this case, firms offer full commitment for all consumers. If:

$$\nu(x_0^*) - cx_0^* \geq \nu_\theta(x_1^*) - cx_1^* - k_1,$$

where $x_1^* \equiv \arg\max_x \nu_\theta(x) - c\,x$, self 1 is not willing to buy any $x > x_0^*$. But he is if:

$$\nu(x_0^*) - cx_0^* \leq \nu(x_1^*) - cx_1^* - k_1.$$

Partially naive consumers will purchase x_0^* at $\tau = 0$, believing that this will be sufficient to overcome self-control problems, but at $\tau = 1$ self 1's preferences have changed more than self 0 anticipated, and it is now optimal for self 1 to go shopping. This means, that partially naive consumer will act like naive consumers, and go shopping. If:

$$\nu(x_0^*) - cx_0^* \leq \nu_\theta(x_\theta^*) - cx_\theta^* - k_1,$$

self 0 will demand a quantity $x_\theta^S \in [x_0^*, x_\theta^*) \quad x_\theta^S < x_0^S$ such that:

$$\nu(x_\theta^S) - c\,x_\theta^S = \nu_\theta(x_\theta^*) - cx_\theta^* - k_1.$$

But per definition of x_0^S, x_θ^S will not be sufficiently big to prevent self 1 from shopping. With partial naïveté, firms will offer the same packaging sizes as in the baseline model, but with introduction of the sizes x_θ^S and $x_1^* - x_\theta^S$. If the consumers in the market have

[1] If $\nu(x) = b(x) - \beta\kappa(x)$ and $u(x) = b(x) - \kappa(x)$ then θ is a measure of the consumer's belief, $\hat{\beta}$, about β. Such that $\theta = 0$ corresponds to $\hat{\beta} = 1$ and $\theta = 1$ corresponds to $\hat{\beta} = \beta$.

continuous levels of naïveté we will have that competitive firms offer a continuous range of packaging sizes.

Circle Model with monopolistic competition

We assume a fraction λ of consumers are naive, and a fraction $1-\lambda$ are sophisticated. Consumers are distributed uniformly along a unit circle. Firms have costs $c\,x + h$ of offering a quantity.

Suppose that n firms have entered the market. In a symmetric equilibrium, firms locate at equal distance along the circle. Consider a consumer located at $\theta \in (0, \frac{1}{n})$ on the circle. Suppose firm i locates at location 0 and firm $i+1$ at location $\frac{1}{n}$. The utility specifications of the consumers in this setting are at time 0 if he buys from firm i $u(x_0) - k_0\left(\frac{1}{2}+\theta\right) - t_0$, and at time 1: $\nu(x_1) - \left(\frac{1}{2}+\theta\right)k_1 - t_1$. The $\left(\frac{1}{2}+\theta\right)$ formulation instead of just θ is to ensure at least partial commitment is possible. If $\theta = 0$ and no other fixed shopping cost exists, a self 0 would never be able to prevent a self 1 consumer from shopping, hence the inclusion of a $\frac{k_0}{2}$. If he buys from firm $i+1$ his utility at time 0 is $u(x_0) - k_0\left(\frac{n+2}{2n} - \theta\right) - t_0$ (analogue for time 1).

Analysis With n firms, making zero profit and x_0^* being sufficient for commitment, the transfers will be:

$$\frac{1}{n}(t_0 - cx_0^*) - h = 0 \quad \Leftrightarrow \quad t = cx_0^* + n\,h.$$

In order to find the level of entry, we need to identify the consumer indifferent between shopping at either of two firms:

$$u(x_0) - k_0\left(\theta + \frac{1}{2}\right) - t' = u(x_0) - k_0\left(\frac{n+2}{2n} - \theta\right) - t''$$
$$\Leftrightarrow \quad \theta_n = \frac{t'' - t'}{2k_0} + \frac{1}{2n}.$$

Since the firm will offer the same transfer to all consumers, the profit maximization problem for pricing an offered variety becomes:

$$\max_{t'} \quad (t' - c\,x_0^*)\,\theta_n.$$

This gives the following first order condition:

$$-\frac{t' - c\,x_0^*}{2k_0} + \left(\frac{t'' - t'}{2k_0} + \frac{1}{2n}\right) = 0.$$

Assuming symmetric strategies $t' = t''$ and we get $t' = t'' = c\,x_0^* + \frac{k_0}{n}$. Using the zero profit condition yields the following level of entry at time 1:

$$\frac{1}{n}\left(c\,x_0^* + \frac{k_0}{n} - c\,x_0^*\right) - h = 0 \quad \Leftrightarrow \quad \left\lfloor \sqrt{\frac{k_0}{h}} \right\rfloor = n_0.$$

In this case, all consumers receive full commitment. If x_0^* is not sufficient for commitment, naive consumers will still demand x_0^* at time 0, and as usual we have, that at time 1, the optimal quantity to offer is $x_1^* - x_0^*$. The difference now being, that naive consumers only constitute a fraction λ of the population. This means that the profit and transfer for a firm serving naive consumers is:

$$\frac{\lambda}{n}(t_0 - cx_0^*) - h = 0 \quad \Leftrightarrow \quad t = cx_0^* + \frac{nh}{\lambda}.$$

Appendix B

The limited Potential of Partial-Commitment

B.0.1 Proposition 2.1

The agent completes the task whenever $-c+\beta b \geq -k \leftrightarrow k+\beta b \geq c$, and gets b but has to pay c. With the residual probability the agent does not complete the task and has to pay k. Therefore, the agent's expected payoff from choosing commitment level k is:

$$U(k) = Pr(-c+\beta b \geq -k)(b - E\left[c|c \leq \beta b + k\right]) - k \cdot Pr(-c+\beta b < -k).$$

And since $Pr(-c+\beta b \geq -k) = F(k+\beta b)$, this becomes:

$$U(k) = F(\beta b + k)(b - E\left[c|c \leq \beta b + k\right]) - k(1 - F(\beta b + k)).$$

To find the optimal interiors solution we can differentiate to get:

$$f(\beta b + k)b - (\beta b + k)f(\beta b + k) - 1 + F(\beta b + k) + k \cdot f(\beta b + k) = 0,$$

where:

$$\frac{\partial}{\partial k}F(\beta b + k)E\left[c|c \leq \beta b + k\right] = (\beta b + k)f(\beta b + k).$$

This means that if an interior best commitment device exists, the following condition holds in optimum:

$$b(1-\beta) = \frac{1 - F(\beta b + k)}{f(\beta b + k)}.$$

For the second order condition to be satisfied it must be that:

$$f'(\beta b+k)b(1-\beta)+f(\beta b+k)<0 \quad \Leftrightarrow \quad b(1-\beta)>\frac{f(\beta b+k)}{-f'(\beta b+k)},$$

where $f'(\cdot)<0$ is necessary. Combined, this gives:

$$b(1-\beta)=\frac{1-F(\beta b+k)}{f(\beta b+k)}>\frac{f(\beta b+k)}{-f'(\beta b+k)}>0.$$

Consider some β' and β'' and some k', then it is always possible to find some k'' s.t. $b\beta'+k'=b\beta''+k'' \Leftrightarrow k''=b(\beta'-\beta'')+k'$. It is therefore possible to vary $b(1-\beta)$ while keeping $\frac{1-F(\beta b+k)}{f(\beta b+k)}>\frac{f(\beta b+k)}{-f'(\beta b+k)}$ fixed. The right hand side, $b(1-\beta)$, is maximized for $\beta=0$ where it takes the value b. Therefore, if:

$$b>\frac{1-F(\beta b+k)}{f(\beta b+k)}>\frac{f(\beta b+k)}{-f'(\beta b+k)}>0,$$

for some $(\beta b+k)\in[\underline{c},\overline{c}]$, then it is possible to find a β s.t. there is an interior optimum, or:

$$b>\frac{1-F(c')}{f(c')}>\frac{f(c')}{-f'(c')}>0,$$

for some $c'\in[\underline{c},\overline{c}]$. The ratio $\frac{1-F(x)}{f(x)}$ is the Mills ratio. Differentiating this gives:

$$\frac{\partial}{\partial x}\frac{1-F(x)}{f(x)}=-\frac{f(x)}{f(x)}-\frac{(1-F(x))}{f(x)^2}f'(x).$$

The derivative is positive if:

$$-\frac{f(x)}{f(x)}-\frac{(1-F(x))}{f(x)^2}f'(x)>0 \Leftrightarrow -f(x)>\frac{(1-F(x))}{f(x)}f'(x),$$

under the assumption that $f'(x)<0$, this gives:

$$-\frac{f(x)}{f'(x)}<\frac{(1-F(x))}{f(x)},$$

which is exactly what follows from the first and second order conditions. Therefore, the necessary conditions to ensure partial commitment is optimal are $f'(c)<0$ and an increasing Mills ratio of $F(c)$. In this case, it will always be possible to find a b large enough for the conditions to be satisfied, but due to the assumption that $\beta b < \underline{c}$, such a b might not overall exist.

B.0.2 Lemma 2.3

For a discrete distribution, the expected payoff from commitment level k is still:

$$U(k) = Pr(-c + \beta b \geq -k)(b - E\left[c|c \leq \beta b + k\right]) - k \cdot Pr(-c + \beta b < -k).$$

Consider the case where there are n cost levels that realize with positive probability. Then the costs can be ordered such that $\underline{c} = c_1 < c_2 < \ldots < c_n = \overline{c}$ with $supp(f) = \{c_1, \ldots ; c_n\}$, and with $f(c_i)$ the probability that c_i realizes for $c_i(f)$. Then the expected payoff from a commitment level k is:

$$U(k) = \sum_{c_i \leq k + \beta b} f(c_i)(b - c_i) - k \cdot \sum_{c_i > k + \beta b} f(c_i),$$

where $c_i \in supp(f)$. For any $k \in [c_i - \beta b, c_{i+1} - \beta b)$ the first sum is constant, but the second sum is increasing in k, therefore, $k = c_i - \beta b$ strictly dominates any $k \in (c_i - \beta b, c_{i+1} - \beta b)$. Therefore, the only commitment devices the agent considers are of the form $k_i = c_i - \beta b$ for $c_i \in supp(f)$. Under the assumption that $b < E[c]$, the agent gets zero payoff from no commitment. This means that the agent will demand partial commitment if there exist a $k_i = c_i - \beta b$ for $c_i \in supp(f)$ s.t. $U(k_i) > 0$, or:

$$\max_{c_i} U(k = c_i - \beta b) = \max_{c_i} \sum_{c=\underline{c}}^{c_i} f(c)(b - c) - k(1 - \sum_{c=\underline{c}}^{c_i} f(c)).$$

Plugging in $k = c_i - \beta b$ and rearranging gives:

$$\max_{c_i} U(k = c_i - \beta b) = \max_{c_i} \beta b - c_i + \sum_{c=\underline{c}}^{c_i} f(c)\left(b(1 - \beta) + c_i - c\right) > 0.$$

B.0.3 Proposition 2.2

If there are only two possible costs $c \in \{\underline{c}, \overline{c}\}$, where $\underline{c} < b < \overline{c}$, and where the probability of $\underline{c}$ occurring is p, partial commitment implies that the time-inconsistent agent completes the task if $c = \underline{c}$ ($\beta b - \underline{c} \geq -k$) but not if $c = \overline{c}$ ($\beta b - \overline{c} < -k$). The agent's preferred partial commitment is therefore $k = \underline{c} - \beta b$.

The agent prefers the partial commitment over no commitment if:

$$(b - \underline{c})p - k(1-p) = (b - \underline{c})p - (\underline{c} - \beta b)(1-p) \geq 0 \Leftrightarrow \beta \geq \frac{\underline{c} - b \cdot p}{b(1-p)}$$

and partial commitment over full commitment if:

$$(b - \underline{c})p - k(1-p)b - \underline{c}p - \overline{c}(1-p) \Leftrightarrow (b - \underline{c})p - (\underline{c} - \beta b)(1-p) \geq b - \underline{c}p - \overline{c}(1-p),$$

which gives:

$$\beta \geq 1 - \frac{\overline{c} - \underline{c}}{b},$$

under the assumption that the task will never be completed absent a commitment device, meaning $\beta b - \underline{c} < 0$. Which gives $\beta < \frac{\underline{c}}{b}$. There always exist ranges of β for which perfect commitment is optimal. For $b < E[c]$:

$$\frac{\underline{c} - b \cdot p}{b(1-p)} < \frac{\underline{c}}{b} \Leftrightarrow b > \underline{c},$$

which is true by assumption. For $b > E[c]$:

$$1 - \frac{\overline{c} - \underline{c}}{b} < \frac{\underline{c}}{b} \Leftrightarrow \overline{c} > b,$$

which is also true by assumption.

B.0.4 Three Costs

Consider an agent with present bias parameter β facing a task with with delayed benefit b and immediate costs c where $c \in \{c_1, c_2, c_3\}$ and where the probability of each cost is

p_1, p_2 and p_3 respectively. Assume $c_1 > \beta b$ — implying absent commitment the task will never be completed — and $c_1 < b < c_3$. The only two partial commitment devices the agent would consider are $k_1 = c_1 - b\beta$ and $k_2 = c_2 - b\beta$. Where k_1 ensures the task will be completed for $c = c_1$ and otherwise not, and k_2 ensures the task will be completed for c_1 and c_2 but not for c_3.

The payoff from choosing k_1 is:

$$U(k_1) = p_1(b - c_1) - k_1(p_2 + p_3) = p_1(b - c_1) - (c_1 - \beta b)(p_2 + p_3).$$

The payoff from choosing k_2 is:

$$U(k_2) = p_1(b - c_1) + p_2(b - c_2) - k_2 p_3 = p_1(b - c_1) + p_2(b - c_2) - (c_2 - \beta b)p_3.$$

The agent prefers k_1 over k_2 if:

$$U(k_1) > U(k_2) \Leftrightarrow \beta > 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}.$$

Choosing k_1 is preferred over no commitment if:

$$U(k_1) = p_1(b - c_1) - (c_1 - \beta b)(p_2 + p_3) > 0 \Leftrightarrow \beta > \frac{c_1 - p_1 b}{p_2 + p_3}.$$

Choosing k_1 is preferred over full commitment if:

$$U(k_1) = p_1(b - c_1) - (c_1 - \beta b)(p_2 + p_3) > b - p_1 c_1 - p_2 c_2 - p_3 c_3 \Leftrightarrow \beta > \frac{b + c_1}{b} - \frac{p_2 c_2 + p_3 c_3}{b(p_2 + p_3)}.$$

Choosing k_2 is preferred over no commitment if:

$$U(k_2) = p_1(b - c_1) + p_2(b - c_2) - (c_2 - \beta b)p_3 > 0 \Leftrightarrow \beta > \frac{c_1 - b(1 - p_3) + (c_2 - c_1)(p_2 + p_3)}{p_3 b}.$$

Choosing k_2 is preferred over full commitment if:

$$U(k_2) = p_1(b - c_1) + p_2(b - c_2) - (c_2 - \beta b)p_3 > b - p_1 c_1 - p_2 c_2 - p_3 c_3 \Leftrightarrow \beta > \frac{c_2 + b - c_3}{b}.$$

These inequalities lead to the following results:

Result B.1. *Given* $b < E[c]$ *and* $c_1 < b < c_2 < c_3$*, a sophisticated agent will:*

Prefer perfect commitment if:

$$\beta > \max\left\{1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}, \frac{c_1 - p_1 b}{p_2 + p_3}\right\}.$$

Complete the task with higher probability than a time-consistent agent if:

$$\beta \in \left\{\frac{c_1 - b(1 - p_3) + (c_2 - c_1)(p_2 + p_3)}{p_3 b}, 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}\right\}.$$

Result B.2. *Given* $b < E[c]$ *and* $c_1 < c_2 < b < c_3$*, a sophisticated agent will:*

Prefer perfect commitment if:

$$\beta \in \left\{\frac{c_1 - b(1 - p_3) + (c_2 - c_1)(p_2 + p_3)}{p_3 b}, 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}\right\}.$$

Complete the task with lower probability than a time-consistent agent if:

$$\beta > \max\left\{1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}, \frac{c_1 - p_1 b}{p_2 + p_3}\right\}.$$

Since neither of these conditions depend on c_3 it is always possible to find a c_3 such that $b < E[c]$ is satisfied. There are several other things to note: Increasing the spread of c but keeping the probabilities fixed increases the likelihood of demanding perfect commitment, and of demanding partial commitment overall. A smaller self-control problem makes the agent more likely to choose the weaker commitment k_1, but it does not necessarily make the agent more likely to behave as a time-consistent agent.

Result B.3. *Given* $b > E[c]$ *and* $c_1 < b < c_2 < c_3$*, a sophisticated agent will:*

Prefer perfect commitment if:

$$\beta > \max\left\{1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}, \frac{b + c_1}{b} - \frac{p_2 c_2 + p_3 c_3}{b(p_2 + p_3)}\right\}.$$

Complete the task with higher probability than a time-consistent agent if:

$$\beta \in \left\{\frac{b + c_2 - c_3}{b}, 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}\right\}.$$

Result B.4. *Given $b > E[c]$ and $c_1 < c_2 < b < c_3$, a sophisticated agent will:*

Prefer perfect commitment if:

$$\beta \in \left\{\frac{b + c_2 - c_3}{b}, 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}\right\}.$$

Complete the task with lower probability than a time-consistent agent if:

$$\beta > \max\left\{1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}, \frac{b + c_1}{b} - \frac{p_2 c_2 + p_3 c_3}{b(p_2 + p_3)}\right\}.$$

Again, if the spread of c increases, keeping the probabilities and expected cost fixed, the probability of an agent preferring partial commitment goes up. For low β, no type of partial commitment is optimal, but as β increases the agent will eventually switch form the stronger commitment to the weaker commitment.

B.0.5 Partial Naivete

The valuation of an intermediate commitment device k is:

$$U(k) = Pr(\beta b - c \geq -k)(b - E[c|\beta b - c \geq -k]) - k \cdot Pr(\beta b - c \geq -k),$$

where:

$$Pr(\beta b - c \geq -k)E[c|\beta b - c \geq -k] = Pr(\beta b - c \geq -k)\int_{\underline{c}}^{\overline{c}} c\frac{Pr(c, \beta b - c \geq -k)}{Pr(\beta b - c \geq -k)}dc,$$

which gives:

$$\int_{\underline{c}}^{\overline{c}} cPr(c)Pr(\beta b - c \geq -k)dc = \int_{\underline{c}}^{\overline{c}} cf(c)\left(1 - \hat{G}\left(\frac{c-k}{b}\right)\right)dc.$$

Define $z = c - \beta b$ and let $F_z(x)$ and $f_z(x)$ be the cdf and pdf of z, then the agent's expected value of an intermediate commitment device k is:

$$U(k) = F_z(k)b - \int_{\underline{c}}^{\overline{c}} cf(c)\left(1 - \hat{G}\left(\frac{c-k}{b}\right)\right) dc - k(1 - F_z(k)).$$

Which leads to the following optimality condition:

$$(k+b)f_z(k) = \frac{1}{b}\int_{\underline{c}}^{\overline{c}} cf(c)\hat{g}\left(\frac{c-k}{b}\right) dc + 1 - F_z(k).$$

B.0.6 Example With Three Cost-Levels

Assume that $b(\overline{\beta} - \underline{\beta}) < c_2 - c_1$. The agent will choose between no commitment, partial commitment with $k_1 = c_1 - \underline{\beta}b$, and partial commitment with $k_2 = c_2 - \underline{\beta}b$. The agent will prefer perfect alignment of preferences over the stricter partial commitment if:

$$U(k_1) = p_1(b - c_1) - k_1(p_2 + p_3) > p_1(b - c_1) + p_2(b - c_2) - k_2 p_3 = U(k_2),$$

which leads to the condition:

$$\underline{\beta} > 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2},$$

which means that for any cost distribution there is a lower bound to $\underline{\beta}$ such that any agent who attaches zero belief to having a worse present bias than $\underline{\beta}$ will prefer a commitment device that perfectly aligns preferences. Assuming that $\overline{\beta} > 1 - \frac{c_2-c_1}{b}\frac{p_2+p_3}{p_2}$, an agent prefers the commitment device with $k = k_1$ over no commitment if $U(k_1) = p_1(b - c_1) - k_1(p_2 + p_3) > 0$, or if:

$$\underline{\beta} > \frac{c_1 - b \cdot p_1}{b - b \cdot p_1}.$$

Therefore, if:

$$\underline{\beta} > \max\{\frac{c_1 - b \cdot p_1}{b - b \cdot p_1}, 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}\},$$

the agent prefers perfect commitment. If:

$$\underline{\beta} < 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2},$$

the agent prefers the stronger commitment level $k = k_2$ over $k = k_1$. The agent prefers k_2 over no commitment if $p_1(b - c_1) + p_2(b - c_2) - k_2 p_3 > 0$, which is the case if:

$$\underline{\beta} > \frac{c_1}{p_3 b} + \frac{c_2 - c_1 - b}{b} \cdot \frac{p_2 + p_3}{p_3}.$$

Overall, the agent prefers stronger than perfect commitment if:

$$\underline{\beta} \in \left(\frac{c_1}{p_3 b} + \frac{c_2 - c_1 - b}{b} \cdot \frac{p_2 + p_3}{p_3}, 1 - \frac{c_2 - c_1}{b}\frac{p_2 + p_3}{p_2}\right).$$

B.0.7 Lemma 2.4

Assume $b(\overline{\beta} - \underline{\beta}) > c_2 - c_1$. In this case, the agent has too much uncertainty for there to be a commitment device that ensures a certain response to each cost realization. No commitment gives zero payoff, and full commitment gives $b - E(c) < 0$. There are three intermediate commitment levels to consider: choosing $k_2 = c_2 - \underline{\beta} b$ ensures that the agent will complete the task with certainty if $c \in \{c_1, c_2\}$ but never if $c = c_3$. This gives the expected payoff:

$$U(k_2) = p_1(b - c_1) + p_2(b - c_2) - k_2 p_3 = p_1(b - c_1) + p_2(b - c_2) - (c_2 - \underline{\beta} b) p_3.$$

But the agent now has to choose between completing the task with certainty if $c = c_1$, by choosing $\overline{k}_1 = c_1 - \underline{\beta} b$, but then having a perceived positive probability of completing the task even for $c = c_2$, which gives the payoff:

$$\begin{aligned} U(\overline{k}_1) &= p_1(b - c_1) + p_2(b - c_2) Pr(b\beta - c_2 > -\overline{k}_1) - \overline{k}_1 \left(p_2(1 - Pr(b\beta - c_2 > -\overline{k}_1) + p_3\right) \\ &= p_1(b - c_1) + p_2(b - c_2)\left(1 - \frac{c_2 - c_1}{b(\overline{\beta} - \underline{\beta})}\right) - (c_1 - \underline{\beta} b)\left(p_2 \frac{c_2 - c_1}{b(\overline{\beta} - \underline{\beta})} + p_3\right), \end{aligned}$$

or ensuring the task is never completed for $c = c_2$, but then having a perceived positive probability of not completing even if $c = c_1$. This can be achieved with the fee $\underline{k}_1 = c_2 - b\overline{\beta}$, and will give the agent an expected payoff of:

$$\begin{aligned} U(\underline{k}_1) &= p_1(b - c_1)Pr(b\beta - c_1 > -\underline{k}_1) - \underline{k}_1(1 - p_1 Pr(b\beta - c_1 > -\underline{k}_1)) \\ &= p_1(b - c_1)\frac{c_2 - c_1}{b(\overline{\beta} - \underline{\beta})} - (c_2 - b\overline{\beta})\left(1 - p_1\frac{c_2 - c_1}{b(\overline{\beta} - \underline{\beta})}\right). \end{aligned}$$

Given the assumptions $c_1 < b < c_2$ and $b(\overline{\beta} - \underline{\beta}) > c_2 - c_1$, it is no longer possible for the agent to construct a commitment device that ensures perfect commitment in expectation. Therefore, the agent expects to complete the task with higher or lower likelihood than a time-consistent agent.

For the case of certain commitment to completing the task for both c_1 and c_2, there is no effect of changing $\overline{\beta}$, but the valuation is increasing in $\underline{\beta}$. This implies that becoming overconfident raises value of strong commitment, if this overconfidence comes through raising the floor on beliefs about β. This effect comes entirely from the reduction in the fee k that is necessary to ensure the desired behavior. For the commitment device insuring the task is always completed for $c = c_1$, increasing $\underline{\beta}$ has ambiguous effects, while increasing $\overline{\beta}$ makes this commitment device strictly more desirable. That is, increasing overconfidence increases the value of this commitment device. Again it is the reduction in fee that dominates. Increasing $\overline{\beta}$ leads the agent to attach a higher probability to the task being completed too often, when $c = c_2$, but this implies incurring the fee $\overline{k}_1$ less often, and the later effect dominates. Finally for the lowest intermediate level of commitment, the value is increasing in $\underline{\beta}$ but ambiguous in $\overline{\beta}$. Increasing $\underline{\beta}$ here has an effect through decreasing the agent's uncertainty, and decreasing the range of β for which the agent would have to forgo completing the task for c_1 to avoid it being completed for c_2. Intuitively, the agent gets closer to the first scenario with complete separation.

B.0.8 Learning about Naivete

There are two possible levels of intermediate commitment that the agent would consider: $\underline{k} = \underline{c} - \overline{\beta}$ and $\overline{k} = \underline{c} - \underline{\beta}$, where the first only ensure the task will be completed for $\underline{c}$ if

$\beta = \overline{\beta}$, and otherwise not. $\overline{k}$ ensures the task will be completed whenever $c = \underline{c}$ and is therefore uninformative. The one-period payoff from picking $\underline{k}$ is:

$$U(\underline{k}) = p(1-g)(b-\underline{c}) - (\underline{c} - \overline{\beta})(1 - p(1-g)) < 0,$$

which is negative by assumption, along with the payoff from picking the higher commitment level:

$$U(\overline{k}) = p(b-\underline{c}) - (\underline{c} - \underline{\beta})(1-p) < 0.$$

These assumptions ensure that a myopic agent would pick no commitment in every period. These two assumptions are equivalent to the following restrictions on β:

$$\underline{\beta} < \frac{\underline{c} - b \cdot p}{b - b \cdot p} \quad \text{and} \quad \overline{\beta} < \frac{\underline{c} - b(1-g)p}{b - b(1-g)p}.$$

For learning to be potentially valuable, knowing ones type has to possibly change the optimal action. Since no commitment is preferred over $\overline{k}$, the agent would choose no commitment if the agent was certain that $\beta = \underline{\beta}$. Therefore, it must be that knowing $\beta = \overline{\beta}$ would lead the agent to choose intermediate commitment.

Intermediate commitment being preferred over no commitment if $\beta = \overline{\beta}$ corresponds to the following condition:

$$U(\underline{k}) = p(b-\underline{c}) - (\underline{c} - \overline{\beta})(1-p) > 0.$$

Therefore, the agent will experiment and choose $k = \underline{k}$ if:

$$p(1-g)(b - \underline{c} + (b - \underline{c})) - (1 - p(1-g))(\underline{c} - \overline{\beta}b + \delta 0) > 0 + \delta 0.$$

Given $\delta = 1$, this means that the agent prefers experimentation if:

$$\overline{\beta} > \frac{\underline{c} - b(1-g)p - (b-\underline{c})p(1-g)}{b - b(1-g)p}.$$

Overall, if:

$$\overline{\beta} \in \left\{ \frac{\underline{c} - b(1-g)p - (b-\underline{c})p(1-g)}{b - b(1-g)p}, \frac{\underline{c} - b(1-g)p}{b - b(1-g)p} \right\},$$

the agent would engage in costly experimentation. To ensure the agent would never complete the task absent a commitment device, $\overline{\beta}b - \underline{c} < 0$ also has to hold. Therefore, $\overline{\beta} < \frac{\underline{c}}{b}$ is also needed. I find that:

$$\frac{\underline{c}}{b} > \frac{\underline{c} - b(1-g)p}{b - b(1-g)p} \quad \Leftrightarrow \quad b > \underline{c},$$

which means that if:

$$\overline{\beta} \in \left\{ \frac{\underline{c} - b(1-g)p - (b-\underline{c})p(1-g)}{b - b(1-g)p}, \frac{\underline{c} - b(1-g)p}{b - b(1-g)p} \right\},$$

and

$$\underline{\beta} < \frac{\underline{c} - b \cdot p}{b - b \cdot p}.$$

Experimentation is considered worthwhile for the agent. If the upper bounds are satisfied but not the lower bound on $\overline{\beta}$, then the agent will not experiment. If g increases (probability of having a severe self-control problem increases), holding everything else fixed, the range of $\overline{\beta}$ for which experimentation is optimal decreases. However, under the assumption that partial commitment is optimal for $\beta = \overline{\beta}$, increasing the spread on the beliefs over present-bias increases the likelihood of experimentation being optimal.

Appendix C

Procrastination and Learning about Self-Control

If $U^I(x;1) > U^I(x;0)$, the time-inconsistent agent will complete the task in periods $t = T - n\underline{\tau}$ where $n \in \mathbb{N}_0$. By Assumption 3.1, the time-inconsistent type will acquire information in period $t = T - 1$ and complete the task in period $t = T$. In period $t = T - 2$, acquiring information means that the task will be completed with perceived probability q in $t = T - 1$ and with $(1 - q)$ in $t = T$. Postponing means information will be acquired with certainty at $t = T - 1$ and completed with certainty at $t = T$. Therefore, the time-inconsistent agent will postpone information acquisition if:

$$\begin{aligned} & -m + qU^I(x;1) + (1-q)U^I(x;2) < -mD(1) + U^I(x;2) \\ \Leftrightarrow \quad & q(U^I(x;2) - U^I(x;1)) > m(1 - D(1)). \end{aligned}$$

Given that this inequality and $\underline{\tau} > 2$ hold, then in $t = T - 3$, acquiring information immediately means the task will be completed in the next period with probability q, but the task will only be completed in period $t = T$ with probability $1 - q$. If the time-inconsistent agent postpones now, she believes that the information will be acquired in the next period with probability q and the task will be completed in two periods from now with probability q, but with probability $1 - q$ the information will be acquired in period $t = T - 1$ and the task will be completed in $t = T$. The time-inconsistent type will therefore postpone information acquisition if:

$$\begin{aligned} & -m + qU^I(x;1) + (1-q)U^I(x;3) \\ & < q(-mD(1) + U^I(x;2)) + (1-q)(-mD(2) + U^I(x;3)) \\ \Leftrightarrow \quad & q(U^I(x;2) - U^I(x;1)) + (1 - qD(1))m > +(1-q)D(3)m. \end{aligned}$$

In general, when $\tau \leq \underline{\tau}$, the time-inconsistent type will postpone acquiring information if

$$q(U^I(x;2) - U^I(x;1)) + (1 - qD(1))m > (1-q)D(\tau - 1)m.$$

The RHS of this inequality is increasing in τ, so if $q(U^I(x;2) - U^I(x;1)) + m(1 - D(1)) > 0$ the time-inconsistent agent will postpone information acquisition for at least $\underline{\tau}$ periods. A sufficient condition for the time-inconsistent agent to procrastinate for $\underline{\tau}$ periods is $U^I(x;2) > U^I(x;1)$.

In $t = T - \underline{\tau} - 1$, the time-inconsistent type knows that if she acquires information, the task will be completed with certainty in the next period, hence her expected payoff is

$$\hat{U}_{T-\underline{\tau}-1}(s_{T-\underline{\tau}-1} = 1) = -m + U^I(x;1).$$

If she chooses not to acquire information, she knows that the time-consistent agent will acquire information in the next period, but the time-inconsistent agent will only acquire information in $t = T - 1$. Hence, her expected payoff is

$$\hat{U}_{T-\underline{\tau}-1}(s_{T-\underline{\tau}-1} = 0) = q(-mD(1) + U^I(x;2)) + (1-q)(-mD(\underline{\tau}) + U^I(x;\underline{\tau}+1)).$$

Hence, she will prefer not acquiring the information in $t = T - \underline{\tau} - 1$ if

$$q(U^I(x;2) - U^I(x;1)) + (1 - qD(1))m > (1-q)(D(\underline{\tau})m + U^I(x;1) - U^I(x;\underline{\tau}+1)).$$

When q approaches 1, this reduces to the previous condition. Whether this condition is generally stronger or weaker than the previous one is not clear. If this assumption holds, however, the agent will procrastinate for an additional $\underline{\tau}$ periods. This follows from the same argument as before.

By induction, for each $t = T - n\underline{\tau}$ such that $t \geq 1$ and $n \in \mathbb{N}_0$, we need to check whether the agent will prefer to acquire information or not. The agent is always comparing acquiring information now and completing the task with certainty tomorrow to the case where the time-inconsistent type will procrastinate until the second to last period and the time-consistent type will acquire information in the next period. Therefore, if

$$q(U^I(x;2) - U^I(x;1)) + (1 - qD(1))m \geq \max_{n,k \in \mathbb{N}} (1-q)(D(n\underline{\tau})m - U^I(x; n\underline{\tau} + k) + U^I(x;k))$$

where $\underline{\tau} = \min\{T, \min\{\tau | U^I(x;\tau) \geq U^I(x;0)\}\}$, then the time-inconsistent agent will procrastinate information acquisition for any $t < T - 1$. Note also that, given this condition, we have derived the time-inconsistent type's learning decision in all continuation games and the equilibrium outcome has been pined down uniquely. □

C.0.1 Derivations of Example 3.1.

For $D(t) = \frac{1}{1+rt}\delta^t$ with $r = \frac{1}{2}$ and $\delta = 1$, $D(t) = \frac{2}{2+t}$. Then Assumption 3.1 is satisfied if $U^I(x;0) = -c + \frac{2}{3}b > 0$, $U^I(x;1) = -\frac{2}{3}c + \frac{2}{4}b > m = 0$, and $U^C(x;0) = -c + b > m = 0$. The strongest of these conditions is $U^I(x;0) = -c + \frac{2}{3}b > 0 \quad \Leftrightarrow \quad \frac{b}{c} > \frac{3}{2}$.

Inequality (3.3) can then be written as:

$$q\left(-\frac{2}{4}c + \frac{2}{5}b + \frac{2}{3}c - \frac{2}{4}b\right) > (1-q)\left(-\frac{2}{3}c + \frac{2}{4}b\right)$$

$$\Leftrightarrow \quad q\left(-\frac{2}{4}c + \frac{2}{5}b\right) > -\frac{2}{3}c + \frac{2}{4}b$$

$$\Leftrightarrow \quad q > \frac{-\frac{2}{3}c + \frac{2}{4}b}{-\frac{2}{4}c + \frac{2}{5}b} = \frac{-20c + 15b}{-15c + 12b}.$$

Note that $-\frac{2}{4}c + \frac{2}{5}b > 0 \quad \Leftrightarrow \quad \frac{b}{c} > \frac{5}{4}$, which is ensured by the assumption of $\frac{b}{c} > \frac{3}{2}$. Plugging $\frac{b}{c} = \frac{23}{15}$ gives $q > \frac{15}{17}$.

For $\frac{m}{c} = \frac{1}{3}$, Assumption 3.1 is satisfied if $U^I(x;0) = -c + \frac{2}{3}b > 0$, $U^I(x;1) = -\frac{2}{3}c + \frac{2}{4}b > m = \frac{c}{3}$, and $U^C(x;0) = -c + b > m = \frac{c}{3}$. The strongest of these conditions is

$U^I(x;1) = -\frac{2}{3}c + \frac{2}{4}b > \frac{c}{3} \quad \Leftrightarrow \quad \frac{b}{c} > 2.$

Then, Inequality 3.3 can then be written as:

$$q\left(-\frac{2}{4}c + \frac{2}{5}b + \frac{2}{3}c - \frac{2}{4}b\right) + \frac{c}{3}\left(1 - q\frac{2}{3}\right) > (1-q)\left(-\frac{2}{3}c + \frac{2}{4}b\right)$$

$$\Leftrightarrow \quad q\left(-\frac{2}{4}c + \frac{2}{5}b - \frac{2}{9}c\right) > -\frac{2}{3}c + \frac{2}{4}b - \frac{1}{3}c$$

$$\Leftrightarrow \quad q > \frac{-c + \frac{2}{4}b}{-\frac{13}{18}c + \frac{2}{5}b} = \frac{-90c + 45b}{-65c + 36b}.$$

Note that $-\frac{13}{18}c + \frac{2}{5}b > 0 \quad \Leftrightarrow \quad \frac{b}{c} > \frac{65}{36}$, which is ensured by the assumption $\frac{b}{c} > 2$. Plugging $\frac{b}{c} = \frac{91}{45}$ gives $q > \frac{5}{39}$. □

C.0.2 Proof of Proposition 3.3.

This follows immediately from the derivations in the main text and comparing the perceived utility of choosing each action. □

C.0.3 Derivations of Examples 3.2 and 3.3.

Time-Consitent Type: Inequality (3.6) can also be written in the following form:

Time-Inconsitent Type: Condition 3.6 can also be written in the following form:

For there to exist parameters b, c such that no information is simultaneously preferred over taking x' and over taking x. The following is required:

$$q > \frac{(b' - m(1+2r))(1+3r)^2}{c\,r^2 + (b' - mr)(1+3r)(1+2r)}.$$

For there to exist parameters b, c such that no information is simultaneously preferred over taking x' and over acquiring information. The following is required:

$$q > \frac{b'(1+r)(1+3r)}{(cr+m(1+3r))(1+2r)}.$$

In our examples $r = \frac{1}{2}$. Therefore, condition C.0.3 reduces to:

$$q > \frac{25b' - 50m}{c + 20b' - 10m}.$$

And condition C.0.3 reduces to:

$$q > \frac{15b'}{4c + 20m}.$$

Overall, from Condition 3.6 and Condition 3.5 we have:

$$1 - \frac{m}{b'} > q > \max\left\{\frac{m}{b-c-b'}, \frac{25b'-50m}{c+20b'-10m}, \frac{15b'}{4c+20m}\right\}$$

Example 3.2: Since $m = 0$ and $\delta = 1$, condition 3.5 is satisfied and condition C.0.3 is binding. Plugging $m = 0$ into C.0.3 directly gives:

$$q > \max\left\{\frac{25b'}{c+20b'}, \frac{15b'}{4c}\right\}.$$

Example 3.3: For $m = \frac{c}{50}$, $b' = \frac{c}{10}$, and $\delta = 1$ condition C.0.3 is satisfied if:

$$\frac{4}{5} > q > \max\left\{\frac{1}{50\frac{b}{c} - 55}, \frac{15}{28}, \frac{15}{44}\right\}.$$

Therefore, if:

$$\frac{4}{5} > \frac{1}{50\frac{b}{c} - 55} \quad \Leftrightarrow \quad \frac{b}{c} > \frac{9}{8}.$$

There will exist a range of q such that the time-consistent agent acquires information immediately, and the time-inconsistent agent will take no action indefinitely. Overall, the conditions for existence of such an equilibrium are: For $\frac{b}{c} \in \left(\frac{9}{8}, \frac{3}{2}\right)$ the time-consistent agent acquires information immediately, and the time-inconsistent agent will take no action indefinitely if:

$$\frac{4}{5} > q > \max\left\{\frac{1}{50\frac{b}{c} - 55}, \frac{15}{28}\right\}.$$

□

C.0.4 Proof of Proposition 3.4.

(i) As in the main text, the assumptions ensure that a time-inconsistent type never works on the task once she learns that her ability is a_L. However, from her $t = 0$ perspective, she wants to complete the task. Also, because $\max_\tau U_t^I(\hat{b};\tau) = U_t^I(\hat{b};1) > 0$ and hence $\max_\tau U_t^C(\hat{b};\tau) = U_t^C(\hat{b};1) > 0$, without learning own ability, both types would immediately learn own self-control and then complete the task. Therefore, such an equilibrium exists. □

(ii) In this case, each type of agent never works on the task once she learns that her ability is a_L. Since information regarding ability is valuable, a time-consistent type chooses to learn in $t = 0$. As $q \to 1$, a time-inconsistent type also (wrongly) believes that, almost surely, she will behave as if she would be time-consistent. Precisely, if Inequality (3.3) holds, then both time-consistent and time-inconsistent types strictly prefer to acquire information about own ability in $t = 0$. Then, the equilibrium behavior for each type in $t \geq 1$ is exactly the same as in Section 3.4. □

Bibliography

Ali, N. S. (2011). Learning Self-Control. *Quarterly Journal of Economics*, 126(2):857–893.

Ariely, D. and Wertenbroch, K. (2002). Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological Science*, 13(3):219–224.

Augenblic, N., Niederle, M., and Sprenger, C. (2013). Working over time: Dynamic inconsistency in real effort tasks.

Bénabou, R. (2015). The economics of motivated beliefs. *Revue d'économie politique*, 125(5):665–685.

Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117(3):871–915.

Bénabou, R. and Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, 112(4):848–886.

Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–164.

Benjamin, D. J., Rabin, M., and Raymond, C. (2016). A model of nonbelief in the law of large numbers. *Journal of the European Economic Association*, 14(2):515–544.

Bernheim, B. D., Meer, J., and Novarro, N. K. (2012). Do consumers exploit precommitment opportunities? evidence from natural experiments involving liquor consumption. Technical report, National Bureau of Economic Research.

Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. (2006). Early decisions: A regulatory framework. Technical report, National Bureau of Economic Research.

Bryan, G., Karlan, D., and Nelson, S. (2010). Commitment devices. *Annual Reviews of Economics*, 2(1):671–698.

Carrillo, J. D. and Mariotti, T. (2000). Strategic ignorance as a self-disciplining device. *Review of Economic Studies*, 67(3):529–544.

Chandon, P. and Wansink, B. (2010). Is food marketing making us fat? a multi-disciplinary review. *Marketing*, 5(3):113–196.

Christensen, E. and Murooka, T. (2018). Procrastination and Learning Self-Control. Working Paper.

Christensen, E. G. B. and Nafziger, J. (2016). Packaging of sin goods – commitment or exploitation? *Journal of Economic Behavior & Organization*, 122:62–74.

Currie, J., DellaVigna, S., Moretti, E., and Pathania, V. (2010). The effect of fast food restaurants on obesity and weight gain. *American Economic Journal: Economic Policy*, 2(3):32–63.

DellaVigna, S. (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature*, 47(2):315–372.

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

DellaVigna, S. and Malmendier, U. (2006). Paying Not to Go to the Gym. *American Economic Review*, 96(3):694–719.

Dobson, P. W. and Gerstner, E. (2010). For a few cents more: Why supersize unhealthy food? *Marketing Science*, 29(4):770–778.

Echenique, F., Imai, T., and Saito, K. (2017). Testable Implications of Models of Intertemporal Choice: Exponential Discounting and Its Generalizations. Working Paper.

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

Eliaz, K. and Spiegler, R. (2008). Consumer Optimism and Price Discrimination. *Theoretical Economics*, 3:459–497.

Esponda, I. and Pouzo, D. (2016). Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models. *Econometrica*, 84(3):1093–1130.

Esteban, S. and Miyagawa, E. (2006). Temptation, self-control, and competitive nonlinear pricing. *Economics Letters*, 90(3):348–355.

Esteban, S., Miyagawa, E., and Shum, M. (2007). Nonlinear pricing with self-control preferences. *Journal of Economic theory*, 135(1):306–338.

Finkelstein, E. A., Trogdon, J. G., Cohen, J. W., and Dietz, W. (2009). Annual medical spending attributable to obesity: payer-and service-specific estimates. *Health affairs*, 28(5):w822–w831.

Fudenberg, D., Romanyuk, G., and Strack, P. (2017). Active learning with a misspecified prior. *Theoretical Economics*, 12(3):1155–1189.

Gagnon-Bartsch, T., Rabin, M., and Schwartzstein, J. (2017). Channeled Attention and Stable Errors. Working Paper.

Galperti, S. (2012). Commitment, flexibility, and optimal screening of time inconsistency.

Gottlieb, D. (2008). Competition over Time-Inconsistent Consumers. *Journal of Public Economic Theory*, 10(4):673–684.

Gottlieb, D. (2014). Imperfect memory and choice under risk. *Games and Economic Behavior*, 85:127–158.

Gottlieb, D. (2016). Will You Never Learn? Self Deception and Biases in Information Processing. Working Paper.

Gul, F. and Pesendorfer, W. (2001). Temptation and Self-Control. *Econometrica*, 69(6):1403–1435.

Heidhues, P. and Kőszegi, B. (2010). Exploiting Naivete about Self-Control in the Credit Market. *American Economic Review*, 100(5):2279–2303.

Heidhues, P., Kőszegi, B., and Strack, P. (2018). Unrealistic Expectations and Misguided Learning. *Econometrica*. Forthcoming.

Hestermann, N. and Le Yaouanq, Y. (2017). It's Not My Fault! Self-Confidence and Experimentation. Working Paper.

Hinnosaar, M. (2012). Time inconsistency and alcohol sales restrictions.

Jain, S. (2012). Marketing of vice goods: A strategic analysis of the package size decision. *Marketing Science*, 31(1):36–51.

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

Kőszegi, B. (2006). Ego Utility, Overconfidence, and Task Choice. *Journal of the European Economic Association*, 4(4):673–707.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(2):443–478.

Laibson, D. (2015). Why don't present-biased agents make commitments? *American Economic Review*, 105(5):267–72.

Lee, H. (2012). The role of local food availability in explaining obesity risk among young school-aged children. *Social science & medicine*, 74(8):1193–1203.

Leung, C. W., Laraia, B. A., Kelly, M., Nickleach, D., Adler, N. E., Kushi, L. H., and Yen, I. H. (2011). The influence of neighborhood food stores on change in young girls' body mass index. *American journal of preventive medicine*, 41(1):43–51.

Murooka, T. and Schwarz, M. A. (2018). The Timing of Choice-Enhancing Policies. *Journal of Public Economics*, 157:27–40.

O'Donoghue, T. and Rabin, M. (1999). Doing It Now or Later. *American Economic Review*, 89(1):103–124.

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

Phelps, E. S. and Pollak, R. A. (1968). On Second-Best National Saving and Game-Equilibrium Growth. *Review of economic Studies*, 35(2):185–199.

Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, pages 141–156.

Schwartzstein, J. (2014). Selective attention and learning. *Journal of the European Economic Association*, 12(6):1423–1452.

Schweighofer-Kodritsch, S. (2017). Time Preferences and Bargaining. *Econometrica*. Forthcoming.

Skiba, P. M. and Tobacman, J. (2008). Payday Loans, Uncertainty and Discounting: Explaining Patterns of Borrowing, Repayment, and Default. Working Paper.

Steenhuis, I. H., Leeuwis, F. H., and Vermeer, W. M. (2010). Small, medium, large or supersize: trends in food portion sizes in the netherlands. *Public health nutrition*, 13(6):852–857.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180.

Vermeer, W. M., Steenhuis, I. H., Leeuwis, F. H., Heymans, M. W., and Seidell, J. C. (2011). Small portion sizes in worksite cafeterias: do they help consumers to reduce their food intake? *International Journal of Obesity*, 35(9):1200.

Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing science*, 17(4):317–337.