Self-rewards and personal motivation

Alexander K. Koch[a] and Julia Nafziger[b*]

[a]Aarhus University and IZA

[b]Aarhus University

June 30, 2009

Abstract

Self-administered rewards are ubiquitous. They serve as incentives for personal accomplishments and are widely recommended as tools for overcoming self-control problems. However, it seems puzzling why self-rewards can work: the prospect of a reward has a motivating force only if the threat of self-denial of the reward after low performance is credible, and the reward is sufficiently enticing. By combining well-received theories of reference-dependent preferences and hyperbolic discounting we explain how a rational forward-looking individual may achieve commitment to self-rewards. Our results show why and when self-regulation built on self-rewards can be successful and thus illustrate the power, but also the limits, of self-rewards.

JEL Classification: A12, C70, D86

Keywords: Self-control, Motivation, Self-reinforcement, Time inconsistency, Reference-dependent preferences

---

*School of Economics and Management, Aarhus University, Building 1322, 8000 Aarhus C, Denmark. [a]Email: akoch@econ.au.dk, Tel.: +45 8942 2137. [b]Email: jnafziger@econ.au.dk, Tel.: +45 8942 2048. We are grateful to Martin Browning, Vincent Crawford, Mathias Dewatripont, Ott Toomet, Frederic Warzynski, and Eyal Winter for helpful comments and discussions.

1 Introduction

The pleasure of the moment often seduces us to act against our long run interests: we are tempted to shirk on unpleasant tasks – such as studying for an exam, writing a report, dieting, or saving money. To increase our motivation we frequently promise ourselves a reward if we persist and do accomplish a particular task. We might, for instance, tell ourselves "If I write three pages for this report, I'll buy a coffee". However, it seems puzzling why such a promise can work without external enforcement: the threat of withholding the reward must be credible, but at the same time the reward must be attractive so that it seems worth the effort. Attempting to motivate myself with a coffee is futile if I know that, should I fail to write anything, I will still buy the coffee.

Nevertheless, self-administered rewards (*self-rewards*) are ubiquitous. For example, the consumer research literature documents the wide-spread use of "self-gifts" as incentives for personal accomplishments in explicit or implicit self-contracts (e.g. Mick and DeMoss 1990, Mick and Faure 1998, Kivetz and Simonson 2002). Moreover, self-rewards are recommended in almost any related self-help guide and figure prominently in the professional treatment of problem behaviors (see e.g. Bandura 1971, 1976, 1986, 2005, Locke, Latham and Erez 1988, Febbraro and Clum 1998, Faber and Vohs 2004, Clum and Watkins 2007). And even firms offer self-management training programs that teach their employees how to use self-rewards (for an overview see e.g. Vancouver and Day 2005).

All this suggests that self-reward strategies must have some power. So why and when is self-regulation built on self-rewards successful? And what determines whether self-rewards can at the same time have motivating force and be credible? The aim of our paper is to shed light on these enduring questions by combining well-received theories of reference-dependent and present-biased preferences.

The need for self-regulation arises from a present-bias (Strotz 1955, Laibson 1997) that creates a self-control problem.[1] From today's perspective the individual finds it optimal to complete a task scheduled in the future. But once the time of the task arrives, she will be tempted to shirk: all else equal, the current cost of effort looms larger (due to the present bias) than the future benefit from task completion. Anticipating her self-control problem, the individual attempts to regulate her future behavior by making a noncommittal promise to herself of a state-contingent self-reward, like buying something nice if (and only if) the task is completed.

Given this promise, buying the good after task completion is something the individual expected to happen; whereas buying it when not completing the task comes unexpectedly. And

[1]There are alternative ways to model intrapersonal conflicts and self-control problems: e.g., Thaler and Shefrin's (1981) *doer-planner model*, Fudenberg and Levine's (2006) *dual self model*, or Gul and Pesendorfer's (2001) *temptation utility model*.

whether expectations are met or not matters: according to Köszegi and Rabin (2006, 2007) past expectations become reference points. When evaluating an outcome relative to such a reference point, people often display loss-aversion in the sense of Kahneman and Tversky (1979).[2] However, the reference point of an individual is not arbitrary: for an individual who is rational and forward-looking it must indeed be optimal (not) to buy the good if she expected before (not) to buy it. That is, actions and expectations must constitute a *personal equilibrium* (Köszegi and Rabin 2006, Köszegi 2009).

To illustrate the consequences of these ideas for self-regulation that this paper develops, consider a self-reward strategy of the form "If I complete the task, I'll buy that nice pair of shoes I saw in the shop window; but if I shirk I will deny myself the shoes." If the price of the shoes is very low, the individual will always buy them – no matter what her past expectations were: the increase in consumption utility outweighs even a psychological loss from spending an unexpected amount of money. Hence, the part of the self-promise "...if I shirk I will deny myself the shoes" is not credible. But then the self-regulation strategy unravels in a personal equilibrium: the individual rationally expects that she will buy the shoes, irrespective of the task outcome. She thus looks forward to the same continuation utility, whether she works on the task or not. This leaves the self-control problem of the individual unresolved, and she will therefore not complete the task. Conversely, if the price of the shoes is very high, she knows that she will never buy them – so the part of the self-promise "If I complete the task, I'll buy that nice pair of shoes..." is not credible. Again this makes it impossible for the individual to rationally expect anything else for the future. In both of these cases, self-rewards are not credible and self-regulation is doomed to fail.

Are promises of self-rewards thus a futile attempt to regulate own behavior? The answer is no, or more precisely, it depends. For an intermediate price range, buying as well as not buying can be part of a personal equilibrium – similar to what Köszegi and Rabin (2006) show: if the individual expects to buy the shoes she will buy them; if she expects not to buy the shoes she will not buy them. The individual can exploit these multiple equilibria for the purposes of self-regulation, because they allow her to believe in the promise that she will buy the shoes upon task completion, but would deny herself the shoes if she failed to work on the task.[3] Thus, the self-reward is credible if its price lies in this intermediate range,

[2]Köszegi and Rabin (2007, p.1048) illustrate this idea as follows: "An employee who had expected a $50,000 salary will assess a salary of $40,000 as a loss, and a taxpayer who had expected to pay $30,000 in taxes will treat a $20,000 tax bill as a gain." Köszegi and Rabin (2006, 2007) discuss evidence that expectations-based counterfactuals influence how people react to outcomes. Evidence for expectations-based reference-dependent preferences comes for example from Crawford and Meng's (2008) reexamination of data on New York city cab driver's labor supply decisions, a lab experiment by Abeler *et al.* (2009), and Mas' (2006) study of police job performance in response to employees' expectations upon entering pay arbitration.

[3]A distinguishing feature from other models where an individual chooses *optimal beliefs* or has *motivated beliefs* (e.g. Akerlof and Dickens 1982, Landier 2000, Brunnermeier and Parker 2005, Gollier and Muermann 2006, Bénabou and Tirole 2007, Brunnermeier, Gollier and Parker 2007) is that beliefs in a

i.e. if the good involved is neither too extravagant nor too much of a bargain or a necessity. Credibility however is not enough. To help overcome the self-control problem the reward must also provide appropriate incentives: completing the task must increase the continuation utility enough to offset the temptation to shirk. This places an additional constraint on the price of the good. It must be sufficiently low, so that buying the good upon task completion is enough of a reward to motivate the individual to put in the required effort.
The more difficult the task is, or the stronger the present bias is, the lower the price of the (given) reward good must be in order to satisfy the incentive constraint, as the comparative statics of our model show. But then this constraint may conflict with the requirement that the self-reward must be credible: namely, if the price of the good is very low, the individual will always buy it, irrespective of her past expectations. In other words, it might happen that whenever the reward is sufficiently powerful to provide incentives, the threat of self-denial is not credible – and whenever self-denial is credible, the reward is not powerful enough. In these cases, self-regulation based on the promise to consume the available reward good only upon task-completion is doomed to fail.
In contrast, self-regulation works if both the credibility and incentive constraints are met. In such a case, the *preferred personal equilibrium* involves a self-reward coupled with rational expectations of the individual that she will complete the task and buy the good upon task completion, but would deny herself the good if she failed to work on the task.

What our arguments sketched above exclude is that the individual can "change her mind" at the point when she faces the effort decision and make herself believe that she will later reward herself even after shirking. This assumption is plausible if the time interval between the task and the self-reward opportunity is short, so that preferences (i.e. the reference point) do not have time to adjust when the individual thinks of a new consumption plan (e.g. "I will buy the shows even if I shirk").
In a model extension based on Köszegi and Rabin (2009) we show that self-regulation still is possible even if the time interval is not short, so that a change of mind does affect preferences. The driving forces of our main analysis remain largely in place: a second lower bound on the price of the reward good ensures that the prospective loss from not originally intended spending on the reward prevents a revision of the self-reward plan. The stronger the individual feels about anticipated changes in future consumption, the lower this price bound, i.e. the easier it is to sustain self-regulation. The additional insight from our model extension is that not only the task difficulty or the severity of the present bias determine the success of self-regulation, but that timely access to the self-reward may also be important: a significant time lag between the effort decision and the self-reward opportunity may allow the individual to readjust her future reference point, making self-regulation more difficult.

personal equilibrium are rational expectations of future outcomes.

Overall, our findings not only offer an explanation for the wide-spread use and success of self-reward strategies, but they also show possible limits to self-regulation. Self-rewards are not a panacea, but work only when adopted in the right circumstances. Our results help to identify these circumstances and thereby provide an explanation for the advice often given in self-help guides. First, rewards should be non-negligible to help motivate yourself, but still "small" so that self-denial after failing to achieve the prescribed target is credible (e.g. "play a game of pinball" rather than "buy the car that I need anyway"). Second, the target should be "appropriate", so that the promised reward has bite (e.g. "write three pages" rather than "write 30 pages"). Third, self-regulation is easier to achieve with "immediate" rewards ("a cup of coffee right after writing" rather than "a cup of coffee sometime later this week").

The paper is organized as follows. After discussing the related literature we introduce the model in Section 2. Our main analysis and results are in Section 3. Section 4 presents a model extension to allow for "changes of mind", and Section 5 concludes the paper.

Related literature

Our main contribution is to the literature that deals with the question of how present-biased individuals cope with self-control problems (for overviews see e.g. Elster 2000, Brocas, Carrillo and Dewatripont 2004). A large body of work focuses on the role of external commitment technologies. It explains why people incur costs – for example by investing in illiquid assets, signing binding contracts, or making binding promises to other parties – in order to overcome self-control problems in savings, consumption and retirement decisions (e.g. Laibson 1997, Diamond and Köszegi 2003, Carrillo 2005), or to overcome low effort provision and procrastination (e.g. O'Donoghue and Rabin 1999b, DellaVigna and Malmendier 2004, Carrillo and Dewatripont 2008).
The growing strand of empirical contributions on this topic shows, for example, that consumers strategically ration their purchase quantities of goods that are likely to be consumed on impulse (Newton *et al.* 1998), and examines the effects of commitment contracts for smokers (Gine, Karlan and Zinman 2008), commitment savings products (Thaler and Benartzi 2004, Ashraf, Karlan and Yin 2006), as well as deadlines (Ariely and Wertenbroch 2002).

While "most instances of self control in everyday life seem to occur without any extrinsic commitment at all", as noted by Rachlin (1995, p.112), only a few papers deal with intrapersonal strategies – as we do here. Most of them assume the presence of an internal commitment device – such as a goal, a deadline, or some control process – and ask how people design and use such devices optimally to regulate behavior (Benhabib and Bisin 2005, Herweg and Müller 2008, Koch and Nafziger 2008).

Our paper endogenously derives conditions under which individuals can achieve internal commitment. It is thus most closely related to Bénabou and Tirole (2004), who study personal rules – another important motivator mentioned in the psychology literature next to self-rewards. In their model, individuals have imperfect recall about past motives, and hence draw inference about these motives based on their past actions (like living up to a personal rule in a situation that puts their willpower to a test). According to Bénabou and Tirole (2004, p.858), a natural interpretation is that in "normal times" the ability to resist impulses "is a known" but it may be hard to predict for "times of stress" (e.g. abstinence). The gain from maintaining a high self-reputation is what can motivate an individual to stick to a personal rule even though the task turned out to be unexpectedly unpleasant (e.g. the individual underestimated how strong the craving for a cigarette or a glass of alcohol is in the morning). Our results complement Bénabou and Tirole (2004) by showing how internal commitment may arise also in routine tasks, for which an individual knows her (in)ability to resist temptation (e.g. writing three pages for a report, or studying for an exam).

2 The model

Overview and timing. We consider the behavior of an individual with present-biased preferences who faces a task. In period 1, the individual chooses whether or not to exert effort to complete the task. The cost of effort is immediate, whereas the outcome (and the utility related to the outcome) realize only in the future. When making the effort decision in period 1, the individual overemphasizes the immediate cost relative to the more distant benefit of task completion, because she has a present bias. This bias creates an intrapersonal conflict of interest: in period 0, the individual thinks that working on the task is optimal. But once she makes the effort choice in period 1, all else equal, she will shirk. Anticipating her self-control problem, the individual can attempt to regulate her behavior by making a noncommittal self-promise in period 0, e.g. of rewarding herself in period 2 if the task was completed.

The task. Completing the task requires effort ($e = \bar{e}$) in period 1 and leads to a delayed benefit $\bar{y}$ at some later date, say in period 2. Should the individual however shirk ($e = \underline{e}$), the task remains unfinished and yields only $\underline{y} < \bar{y}$ in period 2. Effort causes an immediate disutility $c(\bar{e}) = c > 0$, whereas $c(\underline{e}) = 0$.

Self-reward opportunity. There is an opportunity in period 2 to buy one unit of a good at price p. In period 0, the individual can make a noncommittal promise to herself of a consumption plan for period 2 (*self-reward*) that conditions on the period-1 effort decision $e \in \{\underline{e}, \bar{e}\}$. The state-contingent consumption plan $x(e) \in \{0, 1\}$ can for example take the form: "buy the good if the task was completed" ($x(\bar{e}) = 1$) and "don't buy the good if the task remains unfinished" ($x(\underline{e}) = 0$). While for concreteness we will interpret the self-reward as a consumption good that can be purchased at price p, other interpretations are possible: the reward might be any activity that brings pleasure (v) and causes pain or has an opportunity cost (p).

Preferences. When making her effort decision in period 1, the individual overemphasizes the immediate cost relative to the more distant benefit of a completed task. That is, the individual has present-biased preferences (Strotz 1955). Following the literature (e.g. Phelps and Pollak 1968, Laibson 1997, O'Donoghue and Rabin 1999a), we model this using (β, δ)-*preferences* (or *quasi-hyperbolic discounting*). The first parameter, δ, corresponds to the standard exponential discount factor (for simplicity we assume $\delta = 1$). The second parameter, $\beta \in [0, 1)$, captures the extent of the present bias, and is the parameter of interest in our model. The utility of the individual in period $t \in \{0, 1, 2\}$ is given by:

$$U_t = u_t + \beta \left[\sum_{\tau=t+1}^{2} u_\tau \right], \qquad (1)$$

where u_t is the *instantaneous utility* in period t (to be specified below). For instance, the period-0 incarnation of the individual (*self 0*) weighs future utilities u_1 and u_2 equally; but the period-1 incarnation of the individual (*self 1*) puts a larger relative weight on u_1 by discounting u_2 with $\beta < 1$, reflecting her present bias. We assume that the individual is *sophisticated* in the sense of O'Donoghue and Rabin (1999a), i.e. self 0 knows about the present-biased preferences of her future selves.

The instantaneous utility in each period has the properties of Köszegi and Rabin's (2006) reference-dependent utility. It is composed of two components: intrinsic *consumption utility* (which corresponds to the outcome-based utility traditionally studied in economics) and *gain-loss* utility. The latter is related to the deviation of the consumption utility from its reference point (measured in consumption-utility units): losing something that the individual values (e.g. a gold ring) is more painful than losing something that she does not value (e.g. a paper clip). Both consumption utility and gain-loss utility are assumed to be additively separable across the dimensions of the period outcome (see Köszegi and Rabin 2004, 2006 for a discussion of this assumption). We capture these features with a linear loss aversion framework, as in the applications of Köszegi and Rabin (2006) and Heidhues and Köszegi (2008): if an outcome differs in dimension k from its reference point r^k by z, the corresponding gain-loss utility is $\mu^k(z) = \eta^k\, z$ for $z \geq 0$ and $\mu^k(z) = \eta^k\, \lambda\, z$ for $z < 0$. The parameter $\eta^k \geq 0$ measures the weight of gain-loss utility regarding dimension k in the utility function, and $\lambda > 1$ is the *coefficient of loss aversion*.

Specifically, in period 0 no payoff-relevant outcomes occur, and thus $u_0 = 0$. In period 1, the individual makes her effort choice and incurs the cost of effort, which result in a consumption utility of $-c(e)$ and a gain-loss utility of[4] $\mu^e(-c(e) - r^e)$, i.e.

$$u_1 = -c(e) + \mu^e(-c(e) - r^e). \tag{2}$$

In period 2, the individual receives utility from the task outcome $y \in \{\underline{y}, \bar{y}\}$. Furthermore, she has the opportunity to purchase one unit of a good at price $p \geq 0$. Normalizing good and money endowments to (0,0), she derives utility $v(x)$ from $x \in \{0, 1\}$ units of the good and disutility $-p \cdot x$ from spending money:[5]

$$u_2 = v(x) - p \cdot x + y + \mu^x(v(x) - r^x) + \mu^p(-p \cdot x - r^p) + \mu^v(y - r^y), \tag{3}$$

with $v(1) \equiv v > v(0) = 0$.

The reference point as past beliefs, and personal equilibrium. The reference point regarding the outcomes for periods $t > 0$ is determined by the beliefs that the individual forms in period 0. These beliefs are assumed to be rational in the sense of a *personal equilibrium* (Köszegi and Rabin 2006, Köszegi 2009): in each continuation game behavior given beliefs must be consistent with these beliefs. If, for example, self 0 anticipates that she will buy the good in period 2 with probability one, then the reference point of self 2 in the good dimension is $r^x = v$ and in the money dimension it is $r^p = -p$. In a personal equilibrium, it must then indeed be optimal for the individual to buy the good when period 2 is reached, given that her reference point is $(r^x = v, r^p = -p)$.

A crucial feature of this setup is that multiple personal equilibria may exist for some parameter range. For example, in one personal equilibrium the individual expects that she will buy the good, and it is optimal for her to buy the good; in the other personal equilibrium, she expects that she will not buy the good, and then indeed she does not buy the good. When such multiple personal equilibria exist, we assume that self 0 chooses her favorite plan according to Köszegi and Rabin's (2006) notion of a *preferred personal equilibrium*: she implements the personal equilibrium which provides her with the highest expected utility.

Such a plan can include state-contingent consumption, where for example the plan of self 0 to purchase the good only after completing the task translates into corresponding effort-contingent expectations in the good dimension, $\{r^x(\underline{e}) = 0, r^x(\bar{e}) = 1\}$, and in the money dimension $\{r^p(\underline{e}) = 0, r^p(\bar{e}) = -p\}$. To avoid notational clutter, we suppress the effort argument whenever it is unambiguous.

[4] For ease of exposition, we abuse notation and index the dimension for gain-loss utility and the reference point with the respective choice or outcome variable.

[5] We follow Köszegi and Rabin (2006, 2007, 2009) and Heidhues and Köszegi (2008) in assuming that preferences are also over monetary wealth in case the reward is interpreted as buying a good. For a discussion of this assumption see Köszegi and Rabin (2007) or Köszegi and Rabin (2009).

The self-control problem. We will consider the case where the present bias of the individual is sufficiently strong, so that an intrapersonal conflict of interest arises (Assumption 1 below provides the formal condition): in period 0, the individual thinks that working hard to complete the task in period 1 is optimal. But, in the absence of a contingent reward, self 1 will prefer to shirk once she faces the task in period 1.

3 Analysis

Our analysis proceeds as follows. First, we consider the decision of self 2 whether or not to buy the good – given the expectations formed by self 0 (this part of the analysis is based on Köszegi and Rabin 2006). Next, we turn to the effort decision of self 1 at the task stage. Finally, we consider the design of the self-reward scheme by self 0. What expectations can arise in a personal equilibrium given the behavior of self 1 and self 2? And when and how can these be exploited to design a credible and incentive compatible self-reward?

To show more clearly our main insights about designing and committing to self-administered rewards, we make the simplifying assumption that the individual perceives no gain-loss utility related to the task-outcome or effort dimensions, i.e. $\eta^y = \eta^e = 0$. Furthermore, let $\eta^p = \eta^x = 1$. In Appendix A, we allow for the more general case, and show that our results are robust.

3.1 The decision of self 2: to buy or not to buy?

The reference point for self 2 is the individual's past expectation regarding the period-2 outcome. We need to distinguish two cases: did the individual expect to buy the good, or not (in a given state of the world)? Suppose first she did, and her reference points regarding the good and money dimensions thus are $r^x = v$ and $r^p = -p$, respectively. Following through with this expectation provides self 2 with consumption utility $v - p$ from purchasing the good. As her reference points for each outcome dimension are met, she perceives no gain or loss. If, instead, she abstains from buying the good, her consumption utility is 0. So she perceives a loss in the goods dimension and a gain in the money dimension. This results in gain-loss utility $-\lambda\, v + p$. Hence, it is indeed optimal for self 2 to buy the good if $v - p \geq -\lambda\, v + p$, i.e. if

$$p \leq \frac{(1+\lambda)\, v}{2} \equiv p_{MAX}. \tag{4}$$

Should the price exceed p_{MAX}, then self 2 will never consume the good – no matter what expectations the individual held in the past.

Suppose next that the individual did not expect to buy the good, i.e. $r^x = 0$ and $r^p = 0$. Not buying the good in period 2 then provides self 2 with consumption and gain-loss utility of zero. But if she buys the good, her consumption utility is $v - p$ and she perceives gain-loss utility of $v - \lambda p$. Thus, it is optimal for self 2 not to purchase the good if $0 \geq 2v - (1+\lambda)p$, i.e. if

$$p \geq \frac{2v}{1+\lambda} \equiv p_{MIN}. \tag{5}$$

Should the price be less than p_{MIN}, then self 2 will always consume the good – no matter what expectations the individual held in the past.

In the intermediate price range, $p \in [p_{MIN}, p_{MAX}]$, two possibilities arise: if self 0 forms the expectation of buying the good, self 2 will buy the good; and if self 0 forms the expectation of not buying the good, self 2 will not buy the good.

Lemma 1 (Köszegi and Rabin 2006)

1. *Suppose $p < p_{MIN}$. Then self 2 always buys the good, irrespective of past expectations.*
2. *Suppose $p > p_{MAX}$. Then self 2 never buys the good, irrespective of past expectations.*
3. *Suppose $p \in [p_{MIN}, p_{MAX}]$. Then two possibilities exist: self 2 will buy the good if self 0 forms the expectation of buying the good in the given state; and self 2 will not buy good if self 0 forms the expectation of not buying the good in the given state.*

3.2 The decision of self 1: work or shirk?

To understand why and how a contingent reward may act as a positive motivator, we show first that a conflict of interest between self 0 and self 1 may arise in the absence of such a self-regulation strategy. For this, suppose that self 2 will always buy the good, or will never do so, or simply no reward is available in period 2. In other words, only the consumption utility from the task outcome in period 2 varies with the effort that self 1 chooses. What is the effort decision of self 1 in this case? She will work on the task if and only if the gain in consumption utility from task completion arising in period 2, discounted by her present-bias parameter β, exceeds the immediate disutility from effort:

$$\beta(\bar{y} - \underline{y}) - c \geq 0. \tag{6}$$

Self 0, in contrast, does not distort period-1 costs relative to period-2 benefits and thus wants self 1 to complete the task if and only if:

$$(\bar{y} - \underline{y}) - c \geq 0. \tag{7}$$

As $\beta < 1$, Condition (7) may hold while Condition (6) does not. Then an intrapersonal conflict of interest arises: self 1 strictly prefers to shirk, but self 0 would prefer her to put in effort to complete the task.

Assumption 1 (Self-control problem) *There is a conflict of interest between self 0 and self 1, i.e. the individual faces a self-control problem, if and only if:*

$$\frac{c}{\beta} > \bar{y} - \underline{y} \geq c. \tag{SCP}$$

As self-regulation is only an issue if the individual faces a self-control problem, we will henceforth consider the case where Condition (SCP) holds. In an attempt to resolve this intrapersonal conflict of interest, self 0 can make a noncommittal promise of an effort-contingent consumption plan. If the individual will indeed stick to this plan in period 1, then the anticipated continuation utility following task completion is $\bar{y} + \bar{u}_2$, and that if the task remains unfinished is $\underline{y} + \underline{u}_2$. And these distinct continuation utilities can provide incentives for self 1: given that self 0 made a credible promise of a self-administered reward, self 1 completes the task if and only if:

$$\beta\left[(\bar{y} + \bar{u}_2) - (\underline{y} + \underline{u}_2)\right] - c \geq 0 \tag{8}$$

For clarity of exposition we consider the case where the good available in period 2 is a *net reward*, i.e. $v > p$. Thus, a self-reward of the form "If I complete the task I will reward myself and buy the good, if I shirk I will deny myself the good" ($x(\bar{e}) = 1$, $x(\underline{e}) = 0$) will – if the plan is credible – lead to $\bar{u}_2 = v - p > \underline{u}_2 = 0$.[6] The incentive constraint (8) can then be rewritten in terms of a price ceiling p_{IC}:

$$p_{IC} \equiv \bar{y} - \underline{y} + v - \frac{c}{\beta} \geq p. \tag{IC}$$

Put differently, whenever $p_{IC} \geq p$ a self-reward has the power to provide incentives. The more difficult the task is (i.e. the higher c), or the more severe the present bias is (i.e. the lower β), the stronger the incentives must be (i.e. the lower the price ceiling p_{IC}) for the good to still be attractive enough to help overcome the self-control problem.

3.3 Self 0: self-regulation or resignation?

To examine whether a state-contingent consumption plan is not only incentive compatible, but also credible, we need to ask whether it can be part of a personal equilibrium. Three different cases arise as a consequence of Lemma 1. First, if the price of the good is less than p_{MIN}, self 2 will always consume the good – no matter what expectations the individual had. The part of the self-promise "...if I shirk I will deny myself the good" therefore is not credible. In a personal equilibrium, self 0 will always expect her future self to consume the good and self 2 will indeed do so. Hence, $\bar{u}_2 = \underline{u}_2$ and the noncommittal promise of a self-administered reward cannot provide the additional incentives necessary for bridging the conflict of interest between self 0 and self 1. A similar issue arises in the second case, where the price of the good exceeds p_{MAX}. Self 2 will never consume the good – no matter what

[6] Self-reinforcement based on a *net punishment* is analogous. In such a case, the individual specifies $\{x(\bar{e}) = 0,\ x(\underline{e}) = 1\}$, as the consumption benefit v is lower than the disutility p caused along another dimension. While with a self-administered reward self-denial of the reward has to be a credible threat and provide appropriate incentives, with a self-punishment plan the expectation of carrying through with the unpleasant plan in case of shirking has to be credible and provide appropriate incentives.

expectations the individual had. This means that the part of the self-promise "If I complete the task I will reward myself and buy the good..." is not credible. In a personal equilibrium, self 0 will expect not to buy the good and self 2 will indeed not purchase it. Again, $\bar{u}_2 = \underline{u}_2$ and there are no additional incentives for self 1 to exert effort.

In the third case, where $p \in [p_{MIN}, p_{MAX}]$, two personal equilibria arise: when self 0 forms the expectation of buying the good, then it is indeed optimal for self 2 to buy the good; when self 0 forms the expectation of not buying the good, then not buying is indeed optimal for self 2. This multiplicity of personal equilibria provides self 0 with scope to pick her *preferred personal equilibrium*, and thereby motivate her future self with a credible self-administered reward scheme of the form "If I complete the task I will reward myself and buy the good, if I shirk I will deny myself the good." Stated differently, the noncommittal promise is a self-fulfilling plan: given expectations, it is optimal for self 2 to buy the good if the task was completed and not to buy it otherwise. This plan however can only help overcome the self-control problem if the self-reward is sufficiently powerful to motivate self 1, i.e. if $p_{IC} \geq p$ so that the incentive constraint of self 1 holds.

In sum, whenever $p_{IC} \geq p$, in addition to $p \in [p_{MIN}, p_{MAX}]$, the preferred personal equilibrium involves self 0 promising herself a contingent reward plan that motivates self 1 to work on the task. While self 1 is disciplined by the prospect of losing out on the consumption utility from the reward if she does not complete the task, self-denial of the reward does not occur on the equilibrium path and hence does not harm the individual. Hence, forming the expectation not to buy after shirking increases the expected utility of self 0 relative to the alternative without a self-reward plan, where the self-control problem persists and the task does not get completed. We summarize our findings in the following proposition:

Proposition 1

Suppose that self 0 faces a self-control problem, i.e. Condition (SCP) holds. Then for

1. $p < p_{MIN}$ *or* $p > p_{MAX}$, *self-regulation based on a self-administered reward is not credible: in the unique personal equilibrium, self 2 never* ($p > p_{MAX}$) *or always* ($p < p_{MIN}$) *consumes the good and the task does not get completed.*

2. $p_{MAX} \geq p \geq p_{MIN}$, *self-regulation based on a self-administered reward is credible. A good with* $v > p$ *is a sufficiently large reward to provide self 1 with incentives to overcome the self-control problem if and only if in addition* $p_{IC} \geq p$. *Thus, for* $p_{IC} \geq p \geq p_{MIN}$[7] *the preferred personal equilibrium involves self 0 making a promise of buying the good only after task completion and the task gets completed.*

The result sheds light on what on the surface looks like a puzzling phenomenon: self-rewards are widely used and tend to be rather successful when employed in the right circumstances.

[7] Note that $p_{MAX} > v > p_{IC}$, where the last inequality follows from Condition (SCP).

Proposition 1 shows that, under certain conditions, a rational forward-looking individual can indeed overcome a self-control problem with the promise to self-administer a reward after fulfilling a specific target, and the threat to deny herself the reward otherwise.

Two forces may however constrain self-regulation, depending on the characteristics of the available reward good. First, commitment to the self-reward strategy has to be possible, captured by the credibility constraint $p \in [p_{MIN}, p_{MAX}]$. In other words, the good must be neither something that the individual will consume anyway – no matter what her expectations and reference point are (e.g. because it is a necessity or a bargain) – nor so extravagant or expensive that she will never consume it. Second, the reward has to be large enough to motivate the individual to act despite her present bias, captured by the incentive constraint $p \leq p_{IC}$.

These two forces may be in irreconcilable conflict with each other. In order for the good to be a powerful incentive tool it must come at a low enough price ($p \leq p_{IC}$); but then the opportunity to purchase the good may be so attractive that the individual will always end up consuming the good (because $p < p_{MIN}$). The following corollary captures this:

Corollary 1 (Comparative Statics)
The higher the cost of effort c for completing the task is, or the more severe the present bias is (i.e. the lower β), the lower the price ceiling from the incentive constraint $p \leq p_{IC}$. If c is sufficiently high or β is sufficiently low relative to the consumption value v of the good, then p_{IC} lies below p_{MIN}:

$$p_{IC} < p_{MIN} \quad \Leftrightarrow \quad v_{ICMIN} \equiv \frac{1+\lambda}{\lambda - 1}\left[\frac{c}{\beta} - (\bar{y} - \underline{y})\right] > v.$$

In this case, the available good cannot at the same time serve as a reward that provides sufficiently strong incentives for effort and be a credible self-reward.

While credibility of the self-reward does not depend on the task difficulty or the present bias of the individual (c and β do not affect p_{MIN} and p_{MAX}), these factors influence p_{IC}. If the task is very burdensome or the present bias is severe relative to the utility from consuming the good available as a reward, so that $v_{ICMIN} > v$, then self-regulation is doomed to fail: if the reward is sufficiently powerful, the threat of self-denial is not credible – and if self-denial is credible, the reward is not powerful enough.

The result thus helps understand the limits to self-regulation based on self-rewards. The available reward good may come at the "wrong" price, so that either the incentive constraint $p \leq p_{IC}$ or the credibility constraint $p \in [p_{MIN}, p_{MAX}]$ is not met. Or the severity of the self-control problem makes it impossible to satisfy both constraints, regardless of the price at which the available reward good can be purchased, because $v_{ICMIN} > v$. This explains why there still is a large demand for (often costly) external commitment devices.[8] These only

[8]If self-regulation is impossible, self 0 has willingness to pay for external commitment of $\beta\,[\bar{y} - \underline{y} - c]$.

need to satisfy the incentive constraint $p \leq p_{IC}$ to help overcome a self-control problem – credibility is not an issue as the reward is externally enforced.[9]

Our analysis also helps explain the advice that self-help guides give about how to use self-rewards. The comparative statics result above illustrates that it becomes more difficult to find an "appropriate" reward good (if one can choose among several ones) in the face of arduous tasks or when the individual has a strong present bias. Correspondingly, self-help guides typically recommend to start with thinking about suitable contingent rewards: you should use "small rewards" that are sufficiently enticing to help motivate yourself, but that are not so important to you that you would go ahead and get them even if you failed to achieve the prescribed targets (e.g. "play a game of pinball" rather than "buy the car that I need anyway"). What these targets should be is then considered in a second step: if possible, you should divide tasks into relatively unambitious "milestones" that make the cost of reaching a reward small enough relative to the benefit, so that the promised reward for each step has bite (e.g. "write three pages" rather than "write 30 pages").[10]

4 Change of mind?

In our analysis thus far, the individual forms a single set of beliefs in period 0 that determine the reference point to which she later compares outcomes. The model from Section 2 hence captures settings where a "change of mind" in period 1 is without consequences for period 2: self 1 may well think of a new effort-contingent consumption plan, but the reference point does not adjust quickly enough for this to affect what self 2 will do.

Indeed, experimental evidence from Matthey and Dwenger (2008) suggests that reference states of a person do not respond in the short run to changes in beliefs about payoffs. Also, as Tversky and Kahneman (1981, p.456) discuss, bettors at race tracks do not seem to adjust their reference point to accumulated positive or negative earnings from bets. The track takes a percentage of each bet, so many bettors are behind by the end of the day. Studies show that they shift to more risky gambles in late races, in an attempt to avoid a loss for the day (e.g. McGlothlin 1956).

[9] External commitment devices often rely on making it costly to deviate from the preferred action of self 0, i.e. they punish deviations. Schelling (1992) lists examples, such as depositing self-incriminating letters to be published if the individual fails the target. The threatened punishment $v - p < 0$ works if the incentive constraint of self 1 is satisfied, i.e. $[\bar{y} - \underline{y} + p - v] - \frac{c}{\beta} \geq 0$. A self-punishment strategy would – analogously to the self-reward strategy we discussed – require additionally that the threat is credible, i.e. that the individual believes that she would stick to punishing herself (e.g. by engaging in an unpleasurable activity) if she deviated from the prescribed course of action.

[10] For example, if the costs of drafting a report are convex then writing 30 pages in one period is more than ten times as painful as writing three pages in one period. In this case, splitting the task makes the self-control problem for each step less severe in relative terms than the one for the overall task: $\frac{c(3)}{\beta} - \frac{\bar{y}-\underline{y}}{10} < \frac{1}{10}\left[\frac{c(30)}{\beta} - (\bar{y} - \underline{y})\right]$.

But over the medium and long term reference states do seem to adjust, as for instance reflected by the phenomenon of *hedonic adaptation* (see e.g. the survey of Frederick and Loewenstein 2003): a person's reported happiness initially moves strongly following events with lasting consequences – be they negative (e.g. the loss of a limb) or positive (e.g. winning the lottery) – but then tends to revert back to its original mean level.

So what happens in our setting if the time interval between periods 1 and 2 is sufficient for a change of mind by self 1 to result in a change of the reference point for self 2? Does self-regulation unravel? Specifically, will self 1 use the opportunity to revise the consumption plan rather than hold on to the effort-contingent self-reward beliefs formed by self 0?

A change of mind may affect the individual's wellbeing by triggering mental comparisons of the revised beliefs over future consumption with what the individual would have expected to occur according to lagged beliefs: Köszegi and Rabin (2009) and Matthey (2006) discuss that next to *contemporaneous* gain-loss utility (from any contrast between current consumption and the prior beliefs about current consumption) an individual experiences anticipatory or *prospective* gain-loss utility from changes in her beliefs about future consumption. To study the incentives of self 1 to change her mind we therefore apply in the following Köszegi and Rabin's (2009) framework and incorporate prospective gain-loss utility into our model.

4.1 Extending the model

Reference-point adjustment. In contrast to the setting from Section 2, we now allow a change in beliefs to affect the individual's future reference point by assuming, as Köszegi and Rabin (2009) do, that the reference point adjusts *between* periods. Instantaneous adjustment does not seem to be realistic (see our discussion above). In other words, the individual cannot instantaneously move her reference point regarding period-2 consumption (period-1 effort) *when facing* the consumption (effort) decision. But if self 1 forms a new consumption plan rather than follow through with the expectations formed by self 0 (she "changes her mind"), the period-2 reference point regarding consumption adjusts accordingly.

Hence, period 1 is the only relevant opportunity for a change of mind in our three-period model and thus the focus of our model extension. So let us consider more formally what happens in this period. The individual enters period 1 with (possibly effort-contingent) expectations to buy or not too buy the reward good in period 2. These expectations were formed in period 0. Denote the lagged expectations by $r^k_{0,2}(e)$, $k \in \{p, x\}$, $e \in \{\underline{e}, \bar{e}\}$. They serve as the reference point in period 1 for evaluating possible changes in expectations, as described below. During period 1, the individual chooses her effort and simultaneously revises her expectations to $r^k_{1,2}(e)$. If no change occurs, then $r^k_{1,2}(e) = r^k_{0,2}(e)$ for each possible effort level $e \in \{\underline{e}, \bar{e}\}$. These new expectations serve as the reference point for period 2. As before, expectations and associated plans must be rational in the sense that they constitute a personal equilibrium for the continuation game: the individual can make any plan that she knows she will follow through.

Prospective gain-loss utility. Following Köszegi and Rabin (2009), we introduce *prospective* gain-loss utility in the model from Section 2: a change of mind triggers mental comparisons between the revised expectations over future consumption with the lagged expectations over future consumption. The expected utility at the end of period 1 therefore is now composed of three parts. First, consumption utility from outcomes realized within period 1. Second, *contemporaneous* gain-loss utility – experienced because outcomes realized in period 1 are compared with what these outcomes where supposed to be according to the reference point formed in period 0. Third, *prospective* gain-loss utility – which arises because for future dates (period 2) the revised expectations about outcomes for the chosen effort level $e \in \{\underline{e}, \bar{e}\}$, $r^k_{1,2}(e)$, are compared with what was supposed to happen for this effort level according to the lagged expectations, $r^k_{0,2}(e)$. For clarity of exposition, we retain the assumption from Section 3 that the individual experiences no gain-loss utility in the effort and output dimensions (i.e. $\eta^y = \eta^e = 0$).[11] With this simplification, we can write the expected utility of self 1 as

$$u_1 = -c(e) + \underbrace{\gamma_{1,2} \sum_{k \in \{p,x\}} \mu^k \left(r^k_{1,2}(e) - r^k_{0,2}(e) \right)}_{\text{prospective gain-loss utility}}, \tag{9}$$

where, as before, $\mu^k(z) = \eta^k\, z$ for $z \geq 0$ and $\mu^k(z) = \eta^k\, \lambda\, z$ for $z < 0$. As in Section 3 we normalize $\eta^x = \eta^p = 1$. The weight $\gamma_{1,2} \equiv \gamma$ captures the relative importance of changes in beliefs about consumption occurring 1 period in future. The setting of Sections 2 and 3 (as that in Köszegi and Rabin 2006) corresponds to $\gamma = 0$.

4.2 Analysis

The analysis proceeds as follows. First, one has to consider the decision of self 2, whether or not to buy the good. Lemma 1 carries over directly to the model extension, only that self 2 now inherits her reference point from self 1. Next, we turn to the decision of self 1 and whether the expectations formed by self 0 are consistent with the decision of self 1 (and, as in Section 3, the purchase decision of self 2).

A strategy for self 1 consists of an effort-consumption plan combination for each consumption plan that self 0 can possibly form. As our purpose is to examine the robustness of self-regulation, the plan of self 0 that we are interested in is $\{x(\bar{e}) = 1, x(\underline{e}) = 0\}$. This plan leads to self-regulation in the model of Section 2 for $p_{IC} \geq p \geq p_{MIN}$ (see Proposition 1).

[11] As shown in Appendix A, gain-loss utility related to the period-1 effort cost and the period-2 task outcome actually makes the self-control problem less severe. Thus, the case we consider here is the one where the individual has the strongest motivation to revise plans after period 0.

Self-regulation remains possible in the model extension only if self 1 has no incentive to deviate from the effort-plan combination $\{\bar{e}; x(\bar{e}) = 1, x(\underline{e}) = 0\}$. To see this, note first that without a change of plan shirking is not optimal (because $p \leq p_{IC}$). So if self 1 wants to deviate to low effort, she must change the consumption plan. Such a change however implies that self 0 cannot rationally believe in the plan $\{x(\bar{e}) = 1, x(\underline{e}) = 0\}$: after the change of mind by self 1, play in at least one node of the continuation game would contradict her original plan. In a personal equilibrium, self 0 must therefore choose among the alternative plans one that self 1 is willing to carry through with. But $\{x(\bar{e}) = 1, x(\underline{e}) = 0\}$ is the only consumption plan which – if carried through – allows self 0 to overcome the self-control problem, as shown in Section 3.

So when does self 1 have no incentive to deviate? Intuitively (we leave the complete analysis to Appendix B), the most tempting alternative for self 1 is to shirk and make herself believe that she will consume the good after shirking. So let us compare the utility from exerting effort and sticking to the plan of self 0 (which calls for buying the good after high effort) with the expected utility from not exerting effort under a plan that calls for buying the reward good after shirking (unlike the plan formed by self 0). If self 1 exerts effort and expects to buy the good, her expected utility is equal to $\beta\,(\bar{y} + v - p) - c$. Because the period-2 self-reward coincides with what the plan inherited from self 0 says, there is no prospective gain-loss utility. If self 1 shirks and expects to buy the reward good, her expected utility is $\beta\,(\underline{y} + v - p) + \gamma\,(v - \lambda\,p)$. The prospective gain-loss utility in the second part arises because the individual deviates from the no-reward-after-shirking reference point fixed by lagged expectations: the revised consumption plan leads to a prospective gain in period-2 consumption from zero to one unit, and a prospective loss as period-2 expenditure increases from zero to p. Comparing the two expressions, shows that self 1 has no incentive to deviate from her original plan if:

$$p \geq \frac{1}{\gamma\,\lambda}\,\Big[c - \beta\,(\bar{y} - \underline{y}) + \gamma\,v\Big] \equiv p_{CM}. \tag{CM}$$

The condition shows that the price of the reward good must be large enough so that shirking and triggering a change in the consumption plan feels painful, even though self 1 could spare herself the effort cost in the short run.

Proposition 2

Suppose $p_{IC} \geq p \geq p_{MIN}$, i.e. the state-contingent consumption plan that calls for buying the good if the task is completed, and not buying it if the task remains unfinished, is both incentive compatible and credible. Self 1 has no incentive to change her mind, i.e. deviate to any other effort-consumption plan combination if and only if in addition $p \geq p_{CM}$.

Proposition 2 shows that the driving forces of Proposition 1 remain largely intact. What it adds is a second price floor to ensure that the prospective loss from not originally intended spending in period 2 prevents a change of mind by self 1. As long as $p_{CM} \leq p_{MIN}$, this second price floor plays no role: if the price is high enough to ensure that the individual will not want to act against a no-purchase reference point and nevertheless buy the good in period 2, then it is also high enough to prevent herself from changing her mind in period 1. The additional price floor however matters if $p_{CM} > p_{MIN}$. While the range of prices that sustain self-regulation then shrinks, self-regulation still is possible as long as

$$p_{IC} \geq p_{CM} \qquad \Leftrightarrow \qquad v \geq \frac{\beta + \gamma\,\lambda}{\gamma\,(\lambda - 1)} \left[\frac{c}{\beta} - (\bar{y} - \underline{y})\right] \equiv v_{CMIC}.$$

The more important anticipated changes in outcomes (i.e. the larger γ), the smaller v_{CMIC}, and thus the easier it is to find an appropriate reward good. A larger γ means a stronger loss sensation triggered by a change of beliefs away from the reference expectation of not spending money on purchasing the reward good. This loss sensation makes the individual reluctant to give up on plans that she made earlier – even if following through with them leads to costs that she would not incur if she started planning from scratch. The "stickiness" of plans for a sufficiently large γ however is not specific to the self-control problem in our model: Kőszegi and Rabin (2009) and Matthey (2006) obtain similar results for applications where individuals do not have present-biased preferences.

The discussion shows that not only the task difficulty or the severity of the present bias determine the success of self-regulation, but that the time of access to the self-reward may also be important: if there is a significant time span between the actual effort decision and the self-reward opportunity, the individual can readjust her future reference point. As our model extension shows, self-regulation then becomes more difficult because the price of the reward good must exceed an additional lower bound to ensure that the prospective loss from not originally intended spending in period 2 prevents a change of mind by self 1. This suggests that self-rewards work best if they occur promptly after the effort choice is made – something also emphasized in self-help guides (e.g. the University of Victoria Counselling Services 2004) and studies on self-reinforcement (e.g. Grady *et al.* 1988).

5 Conclusion

Three decades ago Bandura raised the question why self-administered rewards can work:

> *"One of the significant, but insufficiently explored, issues in self-reinforcement is why people adhere to contingencies requiring difficult performances, thereby temporarily denying themselves [until requisite performances are attained] rewards over which they exercise full control."* (Bandura 1976, p.140)

We show how a – so far unexplored – combination of well-received theories of reference-dependent preferences and hyperbolic discounting allows to build a tractable model for analyzing self-regulation. Applied to the puzzle of self-rewards, the model explains how commitment to such a strategy can be achieved, and what factors determine whether self-regulation is possible. As we find, self-reward strategies are credible only if the good involved is neither too extravagant nor too much of a bargain or a necessity. But the reward also has to be sufficiently large to provide incentives for the individual to exert effort on the task. Successful self-regulation based on self-administered rewards is possible if the available reward good satisfies these credibility and incentive constraints.

Thus, our results offer an explanation for the wide-spread use and success of self-rewards. But they also help understand the possible limitations of self-rewards: according to the comparative statics of our model, the more difficult the task is, or the stronger the present bias of the individual is, the harder it becomes to jointly satisfy the credibility and incentive constraints. That is, these constraints may conflict with each other: to motivate the individual the reward might have to be so attractive that it becomes impossible for her to believe that she will actually carry through with her threat of self-denial when shirking. For a tough self-control problem it may therefore be difficult to find an "appropriate" reward good.

If there is a significant time span between the effort decision and the self-reward opportunity, the individual can readjust her future reference point by "changing her mind" about the original self-reward strategy. Applying Köszegi and Rabin's (2009) extension of the Köszegi and Rabin (2006) framework, we show that such an opportunity for a change of mind places an additional lower bound on the price of the reward good. This ensures that the prospective loss from not originally intended spending on the reward prevents a change of mind when facing the task. The model extension hence suggests that a longer lag between work on the task and the self-reward opportunity may make commitment to the self-reward more difficult.

Overall, our results provide an explanation for the advice often given in self-help guides. First, rewards should be non-negligible to help motivate yourself, but still "small" so that self-denial after failing to achieve the prescribed target is credible (e.g. "play a game of pinball" rather than "buy the car that I need anyway"). Second, the target should be "appropriate", so that the promised reward has bite (e.g. "write three pages" rather than "write 30 pages"). Third, self-regulation is easier to achieve with "immediate" rewards ("a cup of coffee right after writing" rather than "a cup of coffee sometime later this week").

The limits to self-regulation that our analysis identifies, show why self-rewards are not a panacea. They work only when adopted in the right circumstances, explaining why there still is demand for (often costly) external commitment devices: self-administered rewards can help overcome relatively easy motivational problems (e.g. finishing a report), but the really hard self-control problems (e.g. getting off drugs) require external commitment.

Appendix

A Gain-loss utility for all outcome dimensions

To bring out the main driving forces, we assumed that the individual does not perceive gain-loss utility related to the period-1 effort cost and the period-2 task outcome ($\eta^e = \eta^y = 0$), and we normalized the weight of gain-loss utility in period 2 ($\eta^x = \eta^p = 1$). Here we consider the more general case where, as in Köszegi and Rabin (2006), the individual places equal weight on gain-loss utility across *all* outcome dimensions: $\eta^k = \eta > 0$ for $k = \{e, y, x, p\}$. As before, we assume that Condition (7) holds, i.e. self 0 wants her future self to complete the task.

Now, because self 0 holds expectations regarding the effort of self 1 and the task outcome in period 2, deviations from these reference points trigger gain-loss utility. For instance, with the expectation that the task will be completed, the individual's reference level for the effort cost is $r^e = -c$, and for the task outcome it is $r^y = \bar{y}$. From the perspective of self 1, shirking thus would lead to an immediate gain in the effort-cost dimension of $\eta\, c$ and a discounted loss in the task dimension in period 2 of $-\beta\,\eta\,\lambda\,(\bar{y} - \underline{y})$.

These gain-loss utility components may provide an additional source of motivation for self 1 to complete the task, and therefore change the conditions under which a conflict of interest between self 0 and self 1 arises (in the absence of a self-reward). To see this *task-anticipation effect*, consider the incentives of self 1 in the absence of a self-administered reward. Given the expectation that self 1 will complete the task, it is indeed optimal for self 1 to do so if and only if:

$$\bar{y} - \underline{y} \geq \frac{1+\eta}{1+\eta\,\lambda}\,\frac{c}{\beta} \equiv \Delta\, y'. \tag{10}$$

Effort incentives for self 1 are stronger because the individual suffers a loss if she does not reach the anticipated task outcome.[12] If, on the other hand, self 0 expects self 1 not to provide effort, self 1 will indeed shirk if and only if:

$$\bar{y} - \underline{y} \leq \frac{1+\eta\,\lambda}{1+\eta}\,\frac{c}{\beta} \equiv \Delta\, y''. \tag{11}$$

Because $\Delta\, y'' > \Delta\, y'$, we are left with three cases:

1. $\bar{y} - \underline{y} > \Delta\, y''$: Self 1 will complete the task – no matter what past expectations were.

2. $\bar{y} - \underline{y} < \Delta\, y'$: Self 1 will shirk – no matter what past expectations were.

3. $(\bar{y} - \underline{y}) \in [\Delta\, y', \Delta\, y'']$: Two personal equilibria exist – one with the self-fulfilling expectation of task completion and one with the self-fulfilling expectation of the task

[12] The requirement for the case where $\eta^y = \eta^e = 0$ that $\bar{y} - \underline{y} \geq c/\beta$ is stronger because $\lambda > 1$ implies that $c/\beta > \Delta\, y'$.

remaining unfinished. In the preferred personal equilibrium, self 0 forms the expectation that the task will be completed.

Overall, the condition for a self-control problem to exist hence becomes

$$\Delta y' > \bar{y} - \underline{y} > c. \tag{SCP'}$$

Again, self 0 can attempt to use a self-reward to overcome the self-control problem that arises when Condition (SCP') holds. Analogously to the analysis in the main part, for this to be successful, the self-reward must be credible (i.e. $p \in [p_{MIN}, p_{MAX}]$) and effect of effort on the anticipated period-2 continuation utilities must be sufficient to motivate self 1. The latter means that self 0's expectation that the task will be completed has to be self-fulfilling. This can be part of a personal equilibrium only if the following incentive constraint for self 1 is satisfied:

$$\beta \{(1 + \eta \lambda)(\bar{y} - \underline{y}) + \bar{u}_2 - \underline{u}_2\} \geq (1 + \eta) c. \tag{12}$$

For the case where the good provides net consumption utility $v - p > 0$ we thus obtain a characterization of the preferred personal equilibria that parallels Proposition 1, replacing Condition (SCP) with Condition (SCP') and the threshold p_{IC} with

$$p'_{IC} \equiv (1 + \eta \lambda)(\bar{y} - \underline{y}) + v - \frac{1+\eta}{\beta} c. \tag{13}$$

In sum, the qualitative features of the self-regulation problem are similar to the simplified setting in the main text. What the new conditions show in addition is: i) self-control problems are less likely to arise (but can still arise) if the individual perceives gain-loss utility related to the task outcome and the effort cost dimension; ii) when a self-control problem exists, a higher weight on gain-loss utility η makes it more likely that the individual can overcome it using a self-reward.

B Proof of Proposition 2

Given that $p_{IC} \geq p \geq p_{MIN}$ and that self 0 forms effort-contingent plan $\{x(\bar{e}) = 1, x(\underline{e}) = 0\}$, Part 2 of Proposition 1 ensures that without a change of plan by self 1 effort dominates shirking. What remains to be shown is under what conditions self 1 has no incentive to deviate from the effort-plan combination $\{\bar{e}; x(\bar{e}) = 1, x(\underline{e}) = 0\}$ to any other effort-plan combination.

We do this in two steps. First, we determine which effort level is utility maximizing for each possible alternative plan i) $\{x(\bar{e}) = 0, x(\underline{e}) = 1\}$, ii) $\{x(\bar{e}) = 0, x(\underline{e}) = 0\}$, and iii) $\{x(\bar{e}) = 1, x(\underline{e}) = 1\}$. Second, we compare each of these "maximized" effort-plan combinations to the candidate equilibrium effort-plan combination $\{\bar{e}; x(\bar{e}) = 1, x(\underline{e}) = 0\}$.

1. Utility maximizing efforts for each plan

i) Plan $\{x(\bar{e}) = 1, x(\underline{e}) = 1\}$. If self 1 shirks, her expected utility under this plan is $\beta\,(\underline{y} + v - p) + \gamma\,(v - \lambda\,p)$. If self 1 exerts effort, her expected utility is $\beta\,(\bar{y} + v - p) - c$. Hence, under this plan effort dominates shirking if and only if:

$$p \geq \frac{1}{\gamma\,\lambda}\left[c - \beta\,(\bar{y} - \underline{y}) + \gamma\,v\right] \equiv p_{CM}. \tag{14}$$

ii) Plan $\{x(\bar{e}) = 0, x(\underline{e}) = 1\}$. Under this plan, the expected utility for self 1 from high effort is lower than from shirking: $\beta\,\bar{y} - c + \gamma\,(p - \lambda\,v) < \beta\,(\underline{y} + v - p) + \gamma\,(v - \lambda\,p)$, because $p_{MAX} > v > p_{IC}$ and $c > \beta\,(\bar{y} - \underline{y})$, where the last two inequalities follow from Condition (SCP).

ii) Plan $\{x(\bar{e}) = 0, x(\underline{e}) = 0\}$. Under this plan, the expected utility for self 1 from high effort is lower than from shirking: ii) implies that $\beta\,\bar{y} - c + \gamma\,(p - \lambda\,v) < \beta\,\underline{y}$.

2. Self 1 has no incentive to deviate from $\{\bar{e}; x(\bar{e}) = 1, x(\underline{e}) = 0\}$ to

- $\{\underline{e}; x(\bar{e}) = 0, x(\underline{e}) = 0\}$ if and only if: $\beta\,(\bar{y} + v - p) - c \geq \beta\,\underline{y} \Leftrightarrow p \leq p_{IC}$.
- $\{\underline{e}; x(\bar{e}) = 0, x(\underline{e}) = 1\}$ if and only if: $\beta\,(\bar{y} + v - p) - c \geq \beta\,(\underline{y} + v - p) + \gamma\,(v - \lambda\,p)$ $\Leftrightarrow p \geq p_{CM}$.
- $\{\underline{e}; x(\bar{e}) = 1, x(\underline{e}) = 1\}$ if and only if: $\beta\,(\bar{y} + v - p) - c \geq \beta\,(\underline{y} + v - p) + \gamma\,(v - \lambda\,p)$ $\Leftrightarrow p \geq p_{CM}$.
- $\{\bar{e}; x(\bar{e}) = 1, x(\underline{e}) = 1\}$, as she is indifferent between the two effort-plan combinations.

Thus, self 1 has no incentive to change her mind whenever $p \geq p_{CM}$, in addition to our initial assumption that $p_{IC} \geq p \geq p_{MIN}$.

References

Abeler, Johannes, Armin Falk, Lorenz Götte, and David Huffman, "Reference Points and Effort Provision," IZA Discussion Paper no. 3939 January 2009.

Akerlof, George A. and William T. Dickens, "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, June 1982, *72* (3), 307–319.

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

Ashraf, Nava, Dean Karlan, and Wesley Yin, "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics*, May 2006, *121* (2), 635–672.

Bandura, Albert, "Vicarious and Self-Reinforcement Processes," in Robert Glaser, ed., *The Nature of Reinforcement*, New York: Academic Press, 1971, pp. 228–278.

____, "Self-Reinforcement: Theoretical and Methodological Considerations," *Behaviorism*, 1976, *4* (2), 135–155.

____, *Social Foundations of Thought and Action: A Social Cognitive Theory*, New Jersey: Prentice Hall, 1986.

____, "The Evolution of Social Cognitive Theory," in Ken G. Smith and Michael A. Hitt, eds., *Great Minds in Management: The Process of Theory Development*, Oxford: Oxford University Press, 2005, pp. 9–35.

Bénabou, Roland and Jean Tirole, "Willpower and Personal Rules," *Journal of Political Economy*, August 2004, *112* (4), 848–886.

____ **and** ____, "Identity, Dignity and Taboos: Beliefs as Assets," CEPR Discussion Paper no. 6123 2007.

Benhabib, Jess and Alberto Bisin, "Modeling Internal Commitment Mechanisms and Self-Control: A Neuroeconomics Approach to Consumption-Saving Decisions," *Games and Economic Behavior*, August 2005, *52* (2), 460–492.

Brocas, Isabelle, Juan D. Carrillo, and Mathias Dewatripont, *Commitment Devices under Self-Control Problems: An Overview*, The Psychology of Economic Decisions. Vol.2: Reasons and Choices, Oxford University Press, 2004.

Brunnermeier, Markus K. and Jonathan A. Parker, "Optimal Expectations," *American Economic Review*, September 2005, *95* (4), 1092–1118.

——, **Christian Gollier, and Jonathan A. Parker**, "Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns," *American Economic Review*, May 2007, *97* (2), 159–165.

Carrillo, Juan D., "To be Consumed with Moderation," *European Economic Review*, January 2005, *49* (1), 99–111.

—— **and Mathias Dewatripont**, "Promises, Promises, ...," *Economic Journal*, August 2008, *118* (531), 1453–1473.

Clum, George A. and Patti Lou Watkins, *Handbook of Self-Help Therapies*, New York: Routledge, 2007.

Crawford, Vincent and Juanjuan Meng, "New York City Cabdrivers' Labor Supply Revisited: Reference-Dependence Preferences with Rational-Expectations Targets for Hours and Income," University of California at San Diego, Economics Working Paper No 2008-03. July 2008.

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

Diamond, Peter and Botond Köszegi, "Quasi-Hyperbolic Discounting and Retirement," *Journal of Public Economics*, September 2003, *87* (9-10), 1839–1872.

Elster, Jon, *Ulysses Unbound: Studies in Rationality, Precommitment, and Constraints*, Cambridge, MA: Cambridge University Press, 2000.

Faber, Ronald and Kathleen D. Vohs, "To Buy or Not to Buy? Self-Control and Self-Regulatory Failure in Purchase Behavior," in Roy F. Baumeister and Kathleen D. Vohs, eds., *Handbook of Self-Regulation: Research, Theory, and Applications*, New York, NY: Guilford, 2004, pp. 509–524.

Febbraro, Greg A. R. and George A. Clum, "Meta-Analytic Investigation of the Effectiveness of Self-Regulatory Components in the Treatment of Adult Problem Behaviors," *Clinical Psychology Review*, March 1998, *18* (2), 143–161.

Frederick, Shane and George Loewenstein, "Hedonic Adaptation," in Daniel Kahneman, Edward Diener, and Norbert Schwarz, eds., *Well-Being: The Foundations of Hedonic Psychology*, New York: Russell Sage Foundation, 2003, chapter 16, pp. 302–329.

Fudenberg, Drew and David K. Levine, "A Dual-Self Model of Impulse Control," *American Economic Review*, December 2006, *96* (5), 1449–1476.

Gine, Xavier, Dean Karlan, and Jonathan Zinman, "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation," mimeo, Yale University 2008.

Gollier, Christian and Alexander Muermann, "Optimal Choice and Beliefs with Ex Ante Savoring and Ex Post Disappointment," IDEI Working Paper No. 462 2006.

Grady, Kathleen E., Carol Goodenow, and Joyce R. Borkin, "The effect of reward on compliance with breast self-examination," *Journal of Behavioral Medicine*, 1988, *11* (1), 43–57.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and Self-Control," *Econometrica*, November 2001, *69* (6), 1403–1435.

Heidhues, Paul and Botond Köszegi, "Competition and Price Variation When Consumers Are Loss Averse," *American Economic Review*, September 2008, *98* (4), 1245–1268.

Herweg, Fabian and Daniel Müller, "Performance of Procrastinators: On the Value of Deadlines," Bonn Econ Discussion Papers no. 3-2008 January 2008.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, March 1979, *47* (2), 263–291.

Kivetz, Ran and Itamar Simonson, "Earning the Right to Indulge: Effort as a Determinant of Customer Preferences Toward Frequency Program Rewards," *Journal of Marketing Research*, May 2002, *39* (2), 155–170.

Koch, Alexander K. and Julia Nafziger, "Self-Regulation Through Goal Setting," IZA Discussion Paper no. 3893 December 2008.

Köszegi, Botond, "Utility from Anticipation and Personal Equilibrium," *Economic Theory*, 2009. forthcoming.

___ **and Matthew Rabin**, "A Model of Reference-Dependent Preferences," Economics Department, University of California, Berkeley, Working Paper E04-337 April 2004.

___ **and** ___ , "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, November 2006, *121* (4), 1133–1165.

___ **and** ___ , "Reference-Dependent Risk Attitudes," *American Economic Review*, September 2007, *97* (4), 1047–1073.

___ **and** ___ , "Reference-Dependent Consumption Plans," *American Economic Review*, 2009. forthcoming.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, May 1997, *112* (2), 443–477.

Landier, Augustin, "Wishful Thinking: A Model of Optimal Reality Denial," mimeo, MIT 2000.

Locke, Edwin, Gary P. Latham, and Miriam Erez, "The Determinants of Goal Commitment," *Academy of Management Review*, January 1988, *13* (1), 23–39.

Mas, Alexandre, "Pay, Reference Points, and Police Performance," *Quarterly Journal of Economics*, August 2006, *121* (3), 783–821.

Matthey, Astrid, "Yesterdays Expectation of Tomorrow Determines What You do Today: The Role of Reference-Dependent Utility from Expectations," mimeo, Max-Planck-Institute of Economics Jena October 2006.

___ **and Nadja Dwenger**, "Dont aim too High: The Potential Costs of High Aspirations," mimeo, Max-Planck-Institute Jena October 2008.

McGlothlin, William H., "Stability of Choices among Uncertain Alternatives," *American Journal of Psychology*, 1956, *69* (4), 604–615.

Mick, David Glen and Corinne Faure, "Consumer Self-Gifts in Achievement Contexts: The Role of Outcomes, Attributions, Emotions, and Deservingness," *International Journal of Research in Marketing*, October 1998, *15* (4), 293–307.

___ **and Michelle DeMoss**, "Self-Gifts: Phenomenological Insights from Four Contexts," *Journal of Consumer Research*, December 1990, *17* (3), 322–32.

Newton, Harry N., Cynthia Barnhart, Pamela H. Vance, and Klaus Wertenbroch, "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice," *Marketing Science*, Fall 1998, *17* (4), 317–337.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, March 1999, *89* (1), 103–124.

___ **and** ___ , "Incentives For Procrastinators," *Quarterly Journal of Economics*, August 1999, *114* (3), 769–816.

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

Rachlin, Howard, "Self-control: Beyond commitment," *Behavioral and Brain Sciences*, 1995, *18* (1), 109–159.

Schelling, Thomas C., "Self-Command: A New Discipline," in George Loewenstein and Jon Elster, eds., *Choice Over Time*, New York: Russell Sage Foundation, 1992, chapter 7, pp. 167–176.

Strotz, R. H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 1955, *23* (3), 165–180.

Thaler, Richard H and H M Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy*, April 1981, *89* (2), 392–406.

Thaler, Richard H. and Shlomo Benartzi, "Save More Tomorrow (TM): Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, February 2004, *112* (1), S164–S187.

Tversky, Amos and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science, New Series*, January 1981, *211* (4481), 453–458.

University of Victoria Counselling Services, "Controlling Your Own Study Behaviour," website http://www.coun.uvic.ca/learning/motivation/self-reinforcement.html 2004. last accessed April 2009.

Vancouver, Jeffrey B. and David V. Day, "Industrial and Organisation Research on Self-Regulation: From Constructs to Applications," *Applied Psychology: An International Review*, 2005, *54* (2), 155185.