The Conflict Within Your Brain:

An Economic Treatment of Psychology's Strength Model of Self-Regulation

By

Benjamin Casner

Supervised by

Silvia Sonderegger Ph.D.

This Dissertation is presented in part fulfilment of the requirement for the completion of an MSc. in Behavioural Economics in the School of Economics, University of Nottingham. The work is the sole responsibility of the candidate

Abstract:

There is a strong body of evidence supporting the concept of dual-system or “dual-self” models of self-control. Most theoretical examinations of this concept take the form of principal-agent models with varying mechanisms by which the principal can influence the actions of the agent. However no theoretical examination has incorporated the empirical findings from psychology which support the psychological “strength model” of self-control., most notably that self-regulation is a limited resource which acts like a muscle in that it can be strengthened with exercise and which atrophies with disuse. We develop an economic treatment of the strength model based on findings from Baumeister and other psychologists, then apply it in the context of savings and precommitment, as well as examining the role of initialization costs and framing. Possible extensions to this model include an additional utility cost to self-control, stochastic or uncertain difficulty of self-regulation, and exogenous factors drawing upon the mental energy resource used to self-regulate.

Contents

Last night you, or at least a hypothetical you, stayed up far too late, watched a number of infomercials on TV, and bought a number of useless products that seemed like a good idea at the time. Alternatively, you were working hard reading an academic paper and found yourself eating a combination of pizza and chips for dinner. In another scenario: you signed up for a gym membership, and you go regularly, but only because you want to make sure that you do not waste the money you spent. These are all examples of self-control difficulties or ways to get around them.

The classic economist's model of humanity, Homo economicus, is perfectly rational and perfectly consistent across all time periods, so any decision made at one point in time will be maintained across the whole lifetime of the individual. Homo economicus would be confused by precommitment strategies in single player games where there is no strategic benefit to limiting options. The only time it would need to commit to an action is when doing so yields an advantage in a multi-period game (e.g. disabling a car's steering wheel in a game of chicken).

Schelling (1984) points out that in order for commitment to be rational, there must be two different sets of preferences at work depending on temporal proximity. When an individual yields to temptation, their failure to exert self-control is in essence an act of self-contradiction. The individual does not desire to perform an action since they are aware of the negative future consequences, yet by taking it anyway, they indicate that they also have a significant desire to take the action as well. As McIntosh (1969) put it: "The idea of self-control is paradoxical unless it is assumed that the psyche contains more than one energy system, and that these energy systems have some degree of independence from each other." In order to model these competing systems, Thaler and Shefrin (1981) designed a model based on a principal-agent framework. The principal, referred to as a "planner" makes decisions based on the long term welfare of the individual, whereas the agent "doer" is concerned only with maximizing instantaneous utility. This concept is the basis for most dual-self models of self-control.

The dual-self framework has proven to be both quite flexible and empirically appealing. Brocas and Carrillo (2008) use a dual-self model to explain a number of empirically consistent yet puzzling concepts such as disparity between anticipated future preferences and realized preferences at the time of consumption. They demonstrate that future discounting behavior can be derived within the model itself. Laibson's (1997) quasi-hyperbolic discounting specification could also be considered a simplistic dual-self model. In this model, the future discounting function is given by $\beta\delta^t$, where β represents the individual's present bias and δ is the intertemporal discounting factor. Because the β term has no effect when considering a choice between two consumption bundles in time periods after the present, it can be interpreted as the doer's influence on the actions of the individual.

On a more neuroeconomic note, McClure et al (2004) show that there are two neurological systems which are responsible for decision making. Areas such as the Visual Cortex and Premotor Area activate for all decisions, whereas areas including the Ventral Striatum, Amygdala, and Medial Orbitofrontal Cortex are associated with decisions where rewards are immediately available. Decisions which are only associated with the first areas tend to require effort on the part of the decision maker and are usually based on a chain of logic, whereas decisions involving both systems are more impulsive and tend to be made with little or no thought.

Bechara (2005) refers to these two systems as the reflective and impulsive systems. Individuals in whose brains the reflective areas are damaged or less dominant tend to be more impulsive and prone to addiction. Individuals with impaired impulsive systems can weigh the pros and cons of every option with little difficulty, but often struggle to actually come to a decision. The reflective system exerts control by suppressing instinctive responses on the part of the impulsive system, and editing out distracting sensations during intense cognition. This interaction is why people deep in thought often do not hear things that are said to them, and also why people with little self-control are often easily distracted.

In this paper we examine a variety of empirical and theoretical investigations into the ways which the two decision systems interact, including a summary of a few significant dual-self models in the literature. We also summarize the strength model of self-control from psychology including the work on ego depletion done by Roy Baumeister. We then develop a dual-self economic treatment of the strength model and apply it to explain behavior such as resting after a long period of self-regulation, precommitment, and the effects of framing on self-control decisions. Once we have completed discussion of the applications, we then proceed to consider possible extensions to the model, such as incorporating other factors into the cost function, imperfect knowledge of the cost function (therefore requiring estimation of it), and adding a utility cost to self-control above the ego restrictions.

Self-control as a Limited Resource and the Strength Model of Self-regulation

There are a number of self-control models in psychology, but one of the best supported is the strength model (Baumeister, Vohs, and Tice 2007). The strength model posits that exerting self-control is similar to flexing a muscle, and it has the following properties:

- Self-regulation draws on a central energy supply; when this supply runs low ability to self-regulate is severely diminished until the agent has had a chance to recuperate.
- Like a muscle, frequent exertion of self-control will increase capacity for self-regulation, and neglect will diminish it.

- Because there is a limited store of energy, if anything else draws upon this supply, then capacity to exert self-control will be diminished.

We consider these properties individually and present the evidence for them below:

1. *Self-regulation draws on a central energy supply, and when this supply runs low ability to self-regulate is severely diminished until the agent has had a chance to rest.*

Muraven, Tice, and Baumeister (1998) show that self-control is a limited resource similar to the energy available to a muscle, and that a number of activities draw upon it. Experimental subjects were asked to squeeze a handgrip both before and after viewing an upsetting movie. In the treatment groups, subjects were asked to either heighten or suppress their reaction to the film. The control group was merely asked to view the film with no specific instructions. The treatment group subjects released the handgrip roughly 20 seconds earlier than those in the control group. Similarly, subjects in another experiment who were told to "not think of a white bear" were significantly less persistent when asked to solve an unsolvable anagram puzzle. Baumeister has referred to this resource in other works (Baumeister et al 1998) as "ego" after Freud's theories of mental architecture. We shall adopt this terminology for the remainder of this paper.

Ego levels seem to be dependent on glucose levels in the brain. Gailliot et al (2007) showed that experimental subjects who were supplied milkshakes during a rest period between two tasks requiring willpower performed much better on the later tasks than a control group who were supplied with dull magazines to read. This effect was consistent across both appetizing milkshakes and tasteless concoctions whose appeal was rated even lower than that of the magazines. Ingestion of glucose seems to allow the brain to replenish depleted neurotransmitter reserves at a faster rate than would otherwise be possible.

2. *Like a muscle, frequent exertion of self-control will increase capacity for self-regulation and neglect will diminish it.*

Baumeister, Vohs, and Tice summarized it thus:

> "Targeted efforts to control behavior in one area, such as spending money or exercise, lead to improvements in unrelated areas, such as studying or household chores. And daily exercises in self-control, such as improving posture, altering verbal behavior, and using one's non-dominant hand for simple tasks, gradually produce improvements in self-control as measured by laboratory tasks."
>
> (2007)

These results tend to present as slower ego depletion after a rest period. It is unclear whether this is indicative of greater willpower reserves or more efficient use of the resources available since the experiments in question were not designed to determine the difference.

Benabou and Tirole (2004) propose that there may be an element of self-reputation involved in the changing capacity for self-regulation. Similar to the procrastination model of O'Donoghue and Rabin (1999), given imperfect self-knowledge, an individual deciding whether or not to delay a benefit will estimate the likelihood of being able to resist temptation long enough to receive the rewards of patience based on their past performance. The more successful they have been in the past, the more likely they are to succeed in the future, and so the more likely they will be to try delaying in the first place. If they have a history of being unsuccessful, then they do not consider the expected benefits from delaying worth the cost, and so they will take the smaller immediate benefit. Note that if someone has a history of success based on externally enforced rules, then this is not an accurate measure of ability to self-regulate, and these successes will count for little if anything in estimation of self-control.

The cause for this self-reputation system is most likely due to the fact that information about how the impulsive system values different consumption options may not be consciously available. Bodner and Prelec (2001) point out that many actions are undertaken in order to reinforce some sort of self-image. Someone with little concern for the welfare of the poor may, nevertheless, give a small amount of money to a beggar in order to reinforce a self-perception of altruism. Daniel Kahneman (2010) refers to several different types of utility: anticipatory, feeling, and remembered. This is the utility that occurs before, during, and after consumption respectively, and each type influences behavior in differing ways.

Anticipatory utility is a function of remembered utility, which is not the same as experienced utility. Most people will seek out the consumption options with the highest remembered utility, but this may not be optimal since they may be remembered as being more enjoyable than they actually are. One explanation of how this discrepancy can come about has to do with the way memory works. Few people have the capacity to remember every detail of every event that they experience, so they will usually remember how enjoyable an event is as being somewhere in between the level of the most enjoyable part of it and the average of the event as a whole (Kahneman 2010). For example, an otherwise boring movie with one particularly funny scene may be described as being quite amusing by someone who is not thinking about their experience carefully.

Frederick, Lowenstein, and O'Donoghue (2002) refer to naïve and sophisticated decision makers. A naïve decision maker assumes that preferences are perfectly consistent over time, i.e. that remembered and experienced utility are always identical. A perfectly sophisticated decision maker knows both the remembered and experienced utility functions exactly. Actual decision makers most likely lie somewhere in between. They estimate experienced utility based on various samples of remembered utility.

Putting self-reputation in terms of ego expenditure, someone with a history of failure to self-regulate will have a high estimate for the amount of ego required to perform any given act of self-control, and will, therefore, not attempt it if there is insufficient ego available to spend, even if the amount required is far less than estimated. Commitment based on external rules does not require the individual to flex the self-control "muscle" and so, in addition to providing no information about ability to self-regulate, following these rules will not increase the amount of ego available to spend.

3. Because there is a limited store of energy, if anything else draws upon this supply then capacity to exert self-control will be diminished.

Shiv and Fedorikhin (1999) showed that the contest between the two decision systems can be pushed in favour of one or the other using exogenous factors. They asked subjects to memorize either a two or seven digit number and then proceed to another room and repeat it to an experimenter. On the way they were asked to select one of two dessert options, either chocolate cake or fruit salad. The hypothesis was that the fruit salad, as the healthier option, would be more cognitively appealing, but the chocolate cake would have a greater emotional appeal. Since the subjects who had been asked to memorize a 7 digit number would be under a higher cognitive load, they would be less able to resist temptation and therefore more likely to choose the chocolate cake. This was indeed the case. Subjects under high cognitive load were about 20% more likely to choose the cake than the low load subjects.

Cognitive load is not the only factor which impairs ability to self-regulate. Leith and Baumeister (1996) show that negative emotions increase preference for high risk, high reward actions, even if the expected value of the action is lower than available low risk options. Subjects were presented with a choice of lotteries. The low risk option was a 0.7 probability of winning $2, whereas the high risk option was a 0.04 probability of winning $25. In both cases, losing meant that the subject would be subjected to a loud and unpleasant noise. Participants in the neutral mood condition were told that they would be watching a series of videotapes showing various stimuli as part of a personality test, the experimenter would feign technical difficulties and then present the lottery choice as a way to use the "dead time" effectively. The bad mood condition was similar to the neutral condition except that the subjects were presented with the prospect of singing a highly embarrassing song before the "technical difficulties." Participants in the good mood condition were shown a pair of comedy sketches before being asked to make the lottery choice. Embarrassed subjects were 20% more likely to choose the long shot than those in the neutral condition, and 46% more likely to do so than those in the good mood condition. One possible explanation for this is that bad mood acts as a stressor, which activates a fight or flight response in the individual. Under stressful circumstances, making decisions quickly is often more important than thinking them through carefully, and so the areas of the brain identified by McClure et al as being associated with cognitively intensive decisions may be suppressed.

People are aware of this limitation, at least at a subconscious level. Muraven, Shmueli, and Burkley (2006) performed a series of experiments where they asked subjects to perform a series of self-control tasks. The control group was merely informed of the nature of the tasks, whereas the treatment group was told how many tasks there would be in total. The control group had high performance on the first few tasks, but then fell off quickly. The treatment group's performance on the earlier tasks was not as good as that of the control group, but their curtailed performance in earlier tasks left them with sufficient ego remaining to do better on the later tasks. In a similar experiment, Muraven, Shmueli, and Burkley informed subjects that they would be asked to perform a self-control task for a given period of time, and then surprised them with an additional self-control task at the end of this period. Performance on the surprise task was significantly lower than on the tasks subjects had planned for. This result suggests that people will budget energy for a long series of tasks if they know about them in advance, but will expend their energy quickly if they believe they will only need to exert willpower for a short period of time. This pattern is analogous to the idea that a long distance runner will be much slower than a sprinter but have the endurance to go a great deal farther.

Theoretical Exploration in Previous Literature

The purpose of economics is to understand real world economic situations, but they are far too complex to fully understand with simple observation.

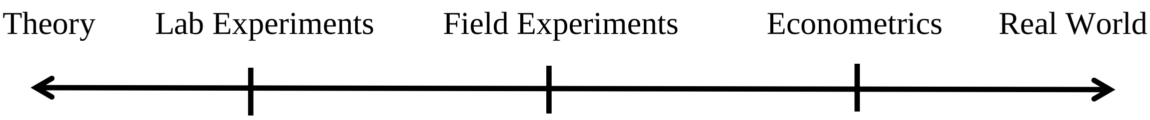

Figure 1

Imagine a line as in figure 1; on one end of this line place theory, and place the real world at the other end. Between these endpoints we have various forms of economic investigation, with the subject of investigation becoming more complex and less controlled as we move from the left to the right. A theorist looking at the real world will formulate a theory of how things work (e.g. people will buy more if prices are lowered). This theory then gets tested at every point along the line, and if it seems to describe the way things occur, it is accepted as an adequate explanation of real world phenomena. If it fails at any point along the line then it is sent back to the theorists to be revised or replaced. This is, perforce, a simplified description of economic methodology, the discussion of what constitutes "failure" alone is contentious at best, but it does provide us with a useful framework for discussing what is required of a good theory.

The purpose of an economic model is, as with a scale model of a vehicle, to create a smaller, less complex version of the thing being modeled in order to get a better sense of the general shape of it and its interactions with the world. A good model will be useful in at least one of the following two functions; the best will fulfil both. First, if reactions to exogenous variables are sufficiently accurate, the model will make predictions that will be useful in real world applications (e.g. using macroeconomic models to design more effective stimulus programs). Second, because models propose a story for the mechanisms behind a phenomenon, they can be tested via the methods further to the right along figure 1. If the theory is deemed valid, then the model has furthered understanding of what it attempts to describe.

The question of what makes for a useful dual-self model is a little easier. As in all economic models, they consist of one or more agents (in this case two) trying to make the most of limited resources. The doer is constrained by wealth/income levels and the self-control actions of the planner. The planner's constraints depend almost entirely on the method of self-control incorporated into the model and the associated cost of self-control. This area is the focus of this paper, and where almost all of the variation in dual-self models originates.

Since the earliest examples, dual-self models have consisted of competition between a myopic and a farsighted self (Frederick, Lowenstein and O'Donoghue 2002). They were somewhat problematic in that which self is in control under which circumstances was poorly defined. Thaler and Shefrin (1981) solved this problem by putting the model into a principal-agent framework, giving control over actions to the doer and endowing the planner with methods to limit the actions of the doer. Their original setup was a fairly simplistic savings model. The method of self-control was to allow the planner to set an upper limit on the instantaneous utility in each period. Consequently, the cost of self-control is exactly the difference between maximal instantaneous utility and restricted utility. Since self-control is accomplished by reducing utility, the constraints on self-regulation are built in. Thaler and Shefrin applied their model in a simple multi-period savings setting where the final period represented retirement. Self-control in each period is set such that marginal disutility of control is equal to the marginal gains in final period utility. By spreading consumption out along the lifetime of the individual, the planner is able to increase total lifetime utility without an increase in income or the savings rate.

Fudenberg and Levine (2006) represent one of the first formal statements of a dual-self model (Frederick, Lowenstein and O'Donoghue 2002). They generalize the Thaler and Shefrin model by starting with a more abstract framework. In each period, the doer chooses an action *a* from the set *A* of all possible actions. This decision is made in the context of starting condition $y \in Y$, and the self-control action *r* chosen by the planner from the set *R* of possible self-regulation options. *A*, *Y*, and *R* are all abstract spaces which Fudenberg and Levine refrain from specifying until they begin work on their applications[1], and they remain entirely agnostic as to the form of R. Instantaneous utility is given by the function $u(y, r, a)$, and the principal's utility is the *discounted* sum of utility in every period over the lifetime of the individual.

The agent's actions are limited by the planner's choice of *R*, but without a cost associated with each value *r*, the principal would have unlimited ability to control the agent. Fudenberg and Levine impose a utility cost to self-control. The cost function to induce an action a is given by:

$$C(y, a) = u(y, 0, a) - u(y, r, a)$$

Equation 1

i.e. the difference between the utility of the restricted action if the agent were to choose it without being induced and the instantaneous utility under self-control. Additionally, they restrict this cost to being strictly greater than zero if $a \notin \arg\max_{a'} u(y, r, a')$. Notice that *y*

[1] The Thaler and Shefrin model is almost a special case of Fudenberg and Levine. Using Thaler and Shefrin as an example, *A* would represent the set of consumption options available to the doer, *Y* represents current wealth levels, and R represents the available range of utility restriction decisions that the planner can choose from.

and a are held constant in the definition of C(.). This cost specification implies that self-control has a utility cost over and above choosing an action that does not maximize instantaneous utility in the given period. This is in contrast to the opportunity cost specification of Thaler and Shefrin or the independent ego resources of the strength model.

Fudenberg and Levine now impose the following assumptions:

a. Any self-control is costly for $r \neq 0$.
b. Self-control is not a finite resource, and therefore for every $a \in A$ there exists an r which induces a.
c. The utility function is continuous.
d. If the agent is indifferent between two actions a, a', then the principal can induce either action at negligible cost.

Using these assumptions they prove that any short-run perfect equilibrium in the principal-agent game is equivalent to the solution of a single player maximization problem, presented here in a somewhat simplified form:

$$max\ U = \sum_{t=1}^{\infty} \delta^{t-1} \big[u(y_t, 0, a_t)\big) - C(y_t, a_t)]$$

Where δ represents the discount rate.

Equation 2

Because the generalized model is agnostic as to the effects of the state variable y on the utility function, Fudenberg and Levine impose a more restrictive assumption that the state interacts with the cost of self-control only in the manner in which it influences the utility of the utility maximizing action and the utility of the induced action (i.e. the opportunity cost). Incorporation of this assumption means that there is a function of the utilities of those two actions which is exactly equal to the cost function. This assumption eliminates the need to specify the form of self-control. In several of their applications, Fudenberg and Levine impose linearity on the cost of self-control[2],

$$C(y,a) = \gamma[\max_{a'} u(y,0,a') - u(y,0,a)]$$

Where $\gamma \geq 0$

Equation 3

The primary problem with these models is that they assume that self-control is an unlimited resource, an assumption which is directly contradicted by Muraven, Tice and Baumeister. Additionally, since Fudenberg and Levine use utility cost as their limiting factor on self-control, they are forced to assume that there is a utility cost to exercising self-control over and above the forgone utility. While this is a fairly intuitive assumption, there is no empirical evidence to support it[3]. It is certainly possible that future experimental work could support utility cost, but the strength model has an advantage in this area because it is, at least in part, derived from experimental evidence. It also has the advantage that utility cost could be introduced into our approach on the strength model fairly easily, but it is more difficult to introduce ego depletion into Fudenberg and Levine or Thaler and Shefrin.

Brocas and Carrillo (2006) adopt a more contract theoretical approach to dual-self modeling. Their focus is not on self-control but rather modeling some findings from neuroeconomics. Their approach is mentioned here because their method to allow the planner to control the actions of the doer is quite unusual in the context of dual-self models, and demonstrates the flexibility of the concept.

The model assumes that the individual is active over a finite number of periods $t \in \{1,2,\dots,T\}$. During each period, the individual undertakes two actions $x_t \in X_T,\ y_t \in Y_T$,

---

[2] If we restrict $\gamma = 1$, and we let *y* be the income stream and *a* the savings/consumption decisions then we get the Thaler and Shefrin savings model. (i.e. Thaler and Shefrin becomes a special case of Fudenberg and Levine with the restriction that $C(y,a) > 0$ removed).

[3] There are confounding factors that could give the appearance of utility cost to self-regulation. For example an agent with reference dependent preferences might have different payoffs when choosing a consumption bundle using self-control than when that bundle is the utility maximizing option.

which can be either pleasant or unpleasant. For example, X could represent labor actions which determine income and Y would then be consumption out of current wealth levels. Given these action spaces, instantaneous utility is given by $U_t(x_t, y_t; \theta_t)$, where θ represents the relative appeal of the two actions. The planner's payoff from the perspective of date *t* is:

$$S_t = \sum_{s=t}^{T} U_s(x_s, y_s; \theta_s)$$

Equation 4

Actions are restricted by an intertemporal constraint

$$B(\{x_s\}_{s=1}^{T}, \{y_s\}_{s=1}^{T}, ; \{\theta_s\}_{s=1}^{T}) \leq 0$$

Equation 5

The general model is agnostic as to the interpretation of B(.). It may represent a budget constraint, an income constraint, or the effects of actions on later payoffs (e.g. the immediate payoff of smoking vs. later health concerns). The planner controls the actions of the doer by offering a series of contracts to the doer subject to equation 5.

While this method is sufficient to model the neurological results mentioned above, it fails to explain how these contracts are enforced. Additionally, this method does not include any form of ego limitation or decay in self-regulatory ability as self-regulation continues. Finally, the idea of a contract based relationship between the two systems seems to be overly anthropomorphizing the interaction of two neurological systems. While context does influence how individuals exercise self-control, it seems to come mainly as a form of budgeting, e.g. "In order to motivate myself to work hard during the week I will indulge myself during the weekend" (Muraven, Shmueli, and Burkley 2006).

Other dual-self models in the literature tend to adopt strategies similar to the ones described above. Notably, Benabou and Pycia (2002) use a strategy similar to that of Thaler and Shefrin or Fudenberg and Levine, but the outcome of the competition between the two systems is determined by a probabilistic function of the resources expended by the planner and doer. Netzer (2009) suggests that the two systems have evolved independent utility functions, and that time inconsistency is a result of prospective consumption switching from one function to the other, similar to the quasi-hyperbolic discounting function. Elster (1985) follows the self-reputation route by proposing that the two selves are in a cartel-like arrangement and that time inconsistency is a result of defection by the doer.

The Model

What we can see from this review of previous work is that very little exploration has focused on the limited nature of self-control. Both Thaler and Shefrin presume that the planner can exhibit unlimited self-control so long as he is willing to pay the utility cost. Brocas and Carrillo do not even place that constraint on the planner's abilities. Any outcome can be induced if a set of contracts can be designed so that it is the most appealing choice. In our modeling process, we emphasize the limited nature of self-control. The planner does not have infinite ego to expend, and must therefore budget self-control expenditure in order to maximize utility.

Being limited by the flow of ego means that the planner will have to take regular rest periods or else have limited capacity for self-regulation. Individuals will be able to save more money towards a specific expenditure, or invest more effort towards a task, if their maximum ego reserve grows or the recovery rate increases. Similarly, this capacity will decrease if costs increase. The effects of increasing capacity to self-regulate will have different implications if the maximum ego reserve grows than if it manifests as reduced ego costs. We show that precommitment to an action leaves more ego available to conserve windfall income than when all self-regulation must be accomplished through willpower alone. Finally, we demonstrate the effects of framing on an individual's ability to complete a task, as well as considering the implications of requiring an ego cost to initiate an activity.

We propose a dual-self model that deviates from the traditional principal-agent framework by giving decision control to the planner rather than limiting the planner to a regulatory role[4]. As in the majority of dual-self models, the planner and doer share the same instantaneous utility function $u(.)$, but the planner's payoff is the sum of the discounted[5] instantaneous utilities from each time period over the lifetime of the individual whereas the doer only cares about instantaneous utility. The planner lives for *n* periods where *n* can be any positive integer up to infinity. If we let U_p and U_d represent the payoffs of the planner and doer respectively:

$$U_p = \sum_{t=1}^{n} \delta^{t-1} u(c_t)$$

Where $0 \leq \delta \leq 1$

Equation 6

and

$$U_d = u(.)$$

Equation 7

The planner has an initial ego endowment $R_1 \geq 0$, with an ego income of *r* per period. R_t represents the planner's ego reserve in period t.

NOTATION:

The maximum possible ego reserve is given by Γ. Therefore $\Gamma \geq R_t \geq 0$ *for all t.*

[4] There is no qualitative difference between this interpretation and a principal expending resources to control the actions of an agent, but we feel that this interpretation is more intuitive.

[5] Discounting is present to ensure that in applications where n is infinite, utility is not infinite for consumption streams where consumption does not eventually fall off to 0. Present bias is not necessary for our purposes since similar behavior can be derived endogenously.

The planner makes a consumption decision c_p from the set *C* of possible consumption bundles in each period with the aim of maximizing long term utility. The doer has the option to override the planner's decision if there is at least one alternative feasible consumption bundle $c_d \in C$ such that $u(c_d) \geq u(c_p)$. The doer will exercise this option in the cases where $u(c_p) \neq argmax_t\, u(c_t)$ (i.e. the cases where it is possible to increase instantaneous utility).

DEFINITION:

Since the doer wishes to maximize instantaneous utility, $c_d = argmax_t\, u(c_t)$*, therefore* $u(c_p) \leq u(c_d)$ *for all* c_p*.*

To this point this model is essentially similar to most dual-self models. Our contribution comes in the self-control method: the planner can expend ego to prevent the doer from exercising the option to overrule the planner's consumption decision.

DEFINITION:

The cost of preventing override by the doer in period t given by

$\omega_t = \omega(.)$*, where* $\omega(.)$ *is an increasing function in the difference between* c_d *and* c_p

$$\omega_t = \omega\left(u(c_d) - u(c_p)\right)$$

Equation 8

PROPERTY 1:

$\omega(.)$ *is increasing in* $u(c_d) - u(c_p)$ *and* $\omega(.) \geq 0\ \forall\ c_d, c_p$

PROPERTY 1a:

$\omega(0) = 0$

Since self-control draws upon an independent resource, it is not present in the planner's utility function. Consequently, the planner suffers no disutility from expending ego and self-regulation is limited only by the ego constraint. Properties 1 and 1a ensure that ego costs are greater than or equal to zero as negative costs would be both nonsensical and unintuitive. Given these properties, we can state the temporal dynamics of the ego reserve in terms of the cost of self-control and the ego income.

$$R_t = R_{t-1} - \omega_{t-1} + r$$

Equation 9

EGO CONSTRAINT:

The planner can choose any c_p such that $\omega_t \leq R_t$. If the planner's desired consumption bundle does not satisfy this constraint then the planner cannot prevent the doer from overriding c_p in period t, however the planner may choose an intermediate consumption level c_r such that $u(c_d) > u(c_r) > u(c_p)$ and $\omega_t = R_t$.

OBSERVATION 1:

The planner's problem is to maximize U_p, therefore in every period t, $c_p \in C$ will be chosen to maximize $\sum_{t=1}^{\infty} \delta^{t-1} u(c_t)$ subject to the ego constraint[6].

[6] This may involve $c_p = 0$ in a number of periods if the planner is conserving resources for later consumption. We see an example of this in the savings application below.

Applications

A. Rest Periods and Labor Investment

One of the most important results of Baumeister's experiment is that ability to exercise self-control is severely diminished once ego resources are depleted. We start with a fairly simple savings model and then work out the boundary condition where the planner prefers to exercise self-control. This model does not compare to Fudenberg and Levine's savings model in that it is not interested in the savings rate. In this application, saving is an "all or nothing proposition". The model says nothing about savings rates in more complex scenarios where it may be optimal to only save part of income.

The application is framed in terms of savings and labor, but it can be used to model a number of behaviors where welfare can be improved by delaying consumption until a sufficient amount of effort has been invested into the endeavour. Some examples could include a student's choice between working on an academic paper and watching a movie, or a dietary choice between quickly prepared meals with poor nutrition and more effortful but healthier meals.

The individual lives for an infinite number of periods ($n = \infty$), in every period the individual can take one of two actions: either labor or consumption. Income from the labor action is denoted by y. Any income earned is placed into a savings pool Y; total savings in period t is denoted Y_t. Savings accumulate interest at rate i. In non-monetary interpretations of this model, the interest rate could be construed as representing the fact that continued effort on a single project can yield benefits greater than focusing on several smaller projects in the same time period.

If the consumption action is taken then the individual will consume some proportion of total savings. We normalize price of consumption to 1, and so $C = \{0, Y_t\}$. Therefore consumption yields instantaneous utility $u(c_t)$ where $c_t \leq Y_t$. If $Y_t = 0$ then the doer will be indifferent between consumption and labor, but if $Y_t > 0$ in any period then $c_d = Y_t$ and so the doer will exercise the option to override the planner's action decision and consume Y_t. The ego cost to take the labor action is:

$$\omega_t = \omega(Y_t - 0) = \omega(Y_t)$$

Equation 10

OBSERVATION 3: *When $Y_t = 0$, $\omega_t = 0$, and so the planner can select the labor action with no ego cost.*

The ego dynamic when the planner is taking the labor action is given by

$$R_t = R_{t-1} - \omega(Y_t) + r$$

Equation 11

ASSUMPTION 1: *In order to ensure limited capacity for self-regulation, we assume r is small so that $\omega(y) > r$. Otherwise the ego constraint could be rendered meaningless.*

Utility with no ego expenditure

If the planner chooses not to expend ego, or if ego is depleted, then the individual will consume all savings, but as noted previously, when $Y_t = 0$ the planner can take the labor action with no interference. Therefore the individual will alternate labor and consumption, and utility will be

$$0 + \delta^2 u(y) + 0 + \delta^4 u(y) + 0 + \delta^6 u(y) + \cdots = \sum_{t=1}^{\infty} \delta^{2t}\, u(y)$$

Equation 12

Since $\delta < 1$ and *y* is constant we can use the formula for the sum of a geometric series to solve this infinite sum:

$$\sum_{t=1}^{\infty} \delta^{2t} u(y) = \sum_{t=1}^{\infty} (\delta^2)^t u(y) = \frac{1}{1-\delta^2} u(y)$$

Equation 13

<u>Utility from a d period savings cycle:</u>

Let *d* be the maximum possible number of periods in which the planner can take the labor action with no rest. The total ego cost will be $\sum_{k=0}^{d} \omega_{t+k}$ and the total ego income over the period will be *dr*. Since *d* is the maximum number of possible periods, initial ego endowment will be equal to the maximum possible reserve Γ. *d* will then be the solution to the equation

$$\sum_{k=1}^{d} \omega_{t+k} = \Gamma + dr$$

Equation 14

We cannot solve for *d* in general since the sum on the left is dependent on the form of the function ω. Fortunately, a general solution is not needed for our results. We can note that *d* will be inversely related to income level *y* and the interest rate *i* given that ω is increasing in *Y*, and both income and interest cause wealth to accumulate more quickly. Similarly *d* will increase with both Γ and *r* since both will increase the ego that the planner has available to expend during the saving cycle.

The rest period to completely refill ego reserves will be $w = \left\lceil \frac{\Gamma}{r} \right\rceil$. The ceiling function represents the fact that if a non-integer number of periods are required to maintain a saving cycle of length *d*, then the individual must rest for the full final period since resting for a fraction of a period is not possible.

After the first period, savings will be equal to the income from that period, i.e. $Y_t = y$. After the second period wealth will be equal to savings, plus accumulated interest and the income from labor in the second period $Y_{t+1} = y(1+i) + y$. Similarly, in the third period we will have $Y_{t+2} = y(1+i)^2 + y(1+i) + y$. Extrapolating to the *d*th period we will then have total wealth equal to $y(1+i)^{d-1} + y(1+i)^{d-2} + \cdots + y$. Using a formula for the sum of a finite geometric sequence we can derive equation 15 which gives savings after *d* periods.

$$y(1+i)^{d-1} + y(1+i)^{d-2} + \cdots + y = \frac{y(1-i^d)}{1-i}$$

Equation 15

During the rest period the individual will not be expending ego and will therefore consume in alternate periods with no ego cost. Utility during the rest period is given by[7]:

$$\delta^{d+2}u(y) + \delta^{d+4}u(y) + \cdots + \delta^{d+w}u(y) = \left[\frac{\delta^{d+2} - \delta^{d+w+2}}{1-\delta^2}\right]u(y)$$

Equation 16

and so utility from the savings cycle will be

$$\sum_{t=1}^{\infty}\left[(\delta^{d+w})^t\left[u\left(\frac{y(1-i^d)}{1-i}\right)\right] + \left[\frac{\delta^{d+2} - \delta^{d+w+2}}{1-\delta^2}\right]^t u(y)\right]$$

Solving for the sum of a geometric series

$$U_p = \frac{u\left(\frac{y(1-i^d)}{1-i}\right)}{1-(\delta^{d+w})} + \frac{u(y)}{1-\left[\frac{\delta^{d+2} - \delta^{d+w+2}}{1-\delta^2}\right]}$$

Equation 17

We do not attempt to solve for the optimal saving cycle here, but we are concerned with the factors which influence the planner to choose saving cycles over alternation. This choice will

[7] Because there are different formulae for odd and even term sums, the equation is written assuming *w* is even to ease use of the formulae, if it is not then subtract 1 from the sum $d + w + 2$ to obtain the correct equation.

only happen if utility from the cycle is greater than that from alternation. With this observation we are ready to state our first proposition:

PROPOSITION 1:

Utility from a savings cycle will be greater than utility from alternation if either of the following conditions are met. Additionally, these conditions are equivalent.

Condition 1:

$$u\left(\frac{y(1-i^d)}{1-i}\right) > u(y)\left(\frac{1}{1-\delta^2} - \frac{1}{1-\theta}\right) * [1-(\delta^{d+w})]$$

$$Where\ \theta = \left[\frac{\delta^{d+2} - \delta^{d+w+2}}{1-\delta^2}\right]$$

Condition 2:

$$u\left(\frac{y(1-i^d)}{1-i}\right) > u(y)\left(\frac{\delta^2\big(1-\delta^2-\delta^d(1-\delta^w)\big)}{(1-\delta^2)\big[1-\delta^2\big(1+\delta^d(1-\delta^w)\big)\big]}\right)[1-(\delta^{d+w})]$$

Using L'Hospital's rule on the middle parenthesis in Condition 2, we find that the term on the right tends to 0 as δ tends to 1, which means that the condition is more likely to be satisfied as the planner becomes more patient. Similarly, going back to Condition 1 and the definition of θ, it is easy to confirm that $\lim_{d\to\infty}\theta = 0$ and $\lim_{w\to\infty}\theta = \frac{\delta^{d+2}}{1-\delta^2}$, which consequently means that the inequality is more likely to be satisfied with an increase in *d* and less likely with an increase in w. In other words, the longer the individual can save, and the fewer periods that need to be spent resting, the more beneficial saving will be. These results are all fairly intuitive given the motivation of the model. Since the savings cycle is an exercise in deferred rewards, increased patience will make the wait for those rewards more bearable. An increase in *d* will increase the amount the individual will be able to save before giving in to temptation, whereas an increase in w represents time when the individual is too tired to save and is thus forced to alternate labor and consumption.

The effects of an increase in income or interest are ambiguous. Both will reduce d, but they will also increase the utility term on the left side of the conditions. Therefore the effect depends on whether the increase in benefit to saving outweighs the decreased capacity to save, which is in turn dependent on the specification of ω. An increase in i will reduce the amount of time the planner can save, but increases the utility from saving without affecting the utility from alternation. Therefore interest will be more likely to increase the appeal of saving than to decrease it. An increase in y will increase the utility term on the right as well as the term on the left, though the term on the left will increase more since saving amplifies the effect of raised wages. Since the temptation to consume immediately will also increase with an increase in income, it is impossible to say whether an increase in income will increase the likelihood of saving without specifying ω.

d and w both increase with the maximum ego reserve Γ, therefore the effects of an increase in total ego potential on utility depend on the relative size of the changes in the two variables, which in turn depend on the values of ω and r. An increase in ω will reduce d, and therefore reduce utility from savings. r increases d and reduces w, so an increase in ego income will increase savings cycle utility.

We can model increasing capacity for self-regulation as either an increase in Γ or a decrease in ω. Both changes will increase d, but decreasing costs will have little effect on w. In other words, an increase in maximum ego capacity will increase the required waiting time to get the benefits of increased capacity for self-regulation, but decreasing the cost of self-control will not. In either case, the increase in self-control over time must be at least asymptotically limited, similar to the way human muscle strength seems to be limited, in order to prevent the savings period from becoming effectively infinite.

One interesting point to note is that the model explains some of the evolutionary function of the impulsive systems. In an extreme case, without the presence of the doer to actually force the individual to consume, the planner could keep saving throughout the entire lifetime of the individual and never actually consume anything. While most researchers focus on the need to mitigate the effects of the doer, this model shows that *completely eliminating its influence would be almost as deleterious as giving it complete control*. One possible avenue for future research would be to see whether individuals who exhibit excessive savings behavior such as compulsive hoarders or people who save excessive amounts "for a rainy day" are less impulsive or if the impulsive systems in their brains are less active than normal.

B. Precommitment, Distractions, and Conservation of Windfall Income

In this section, we examine the effects of precommitment on an individual's ability to save windfall income as well as the effects of precommitment on saving out of predictable income flows. Since we are concerned with ability to save rather than utility, we will set $\delta = 1$ and restrict $u(.)$ to a linear form. These restrictions ensure that the planner's incentive is to maximize saving in order to consume everything in one large consumption binge in the (infinitely distant) final period. In our application we assume that the individual's lifetime is infinite, but we can interpret this assumption as a lifetime longer than the planner can actually conceive of, and so lifetime is treated as infinite by default. Saving would then be saving for retirement which is regarded as a single period rather than a portion of the individual's life.

The individual has the same consumption and labor options as before, but there is an additional action where the individual can place wealth into an illiquid savings asset. Withdrawal of wealth from this savings asset must be initiated in the previous period, so the doer has no incentives to try to access wealth stored within it.

At the end of each period the individual has a probability p of receiving windfall income S. S is large enough such that $\omega(S) = \Gamma$. We know from the previous section that the planner's incentive is to save as much as possible, therefore the planner will expend all available ego and save a proportion $s < S$ of the windfall such that:

$$\omega(s + Y_t) = R_t$$

Equation 18

If the planner makes use of the illiquid saving option, then the dynamics are significantly different than in the previous section. In the first period, the individual has no wealth and will therefore take the labor option. Since the doer cannot withdraw savings from the illiquid asset, there is no ego cost associated with keeping savings in that account. Therefore the planner will want to place income from the labor action into the illiquid savings asset in the second period at ego cost $\omega(y)$. In the third period the individual again has no liquid wealth and therefore can take the labor action at no ego cost. The dynamic after this point is dependent on $\omega(y)$ and r. The ego dynamic after the first two periods will be given by

$$R_{t+2} = R_t + r + r - \omega(y)$$

$$\Rightarrow R_{t+2} = R_t + 2r - \omega(y)$$

Equation 19

From equation 19 we can see that the individual will alternate labor and placing earned income in the illiquid saving asset indefinitely so long as the inequality below is satisfied

$$2r \geq \omega(y)$$

Equation 20

If equation 20 is not satisfied then the planner has two options. Strategy *S1* is to take a brief rest period by alternating labor and consumption for a short while after saving for a few periods.

Strategy *S2* is to continue placing savings into the saving asset until ego is depleted and then engage in a longer rest period until it is replenished. Since the planner wishes to maximize saving, they will choose the option which has the least total rest time. As discussed in the previous section, the long waiting period will be $w = \left\lceil \frac{\Gamma}{r} \right\rceil$ and it will be necessitated after *d* periods with *d* determined by equation 14. Ego reserves after saving for *n* periods will be $\Gamma - \sum \omega_t + nr$, so subtracting this from Γ we get a rest period of $x = \left\lceil \frac{\sum \omega_t - nr}{r} \right\rceil$ periods. The proportion of waiting periods to saving periods in each situation is $\frac{w}{d}$ and $\frac{x}{n}$ respectively. The strategy which has the least waiting time is the strategy whose representative ratio is the lowest. This proves proposition 2

PROPOSITION 2: *If the planner has an illiquid savings asset available, then the planner will prefer strategy S2 if and only if the following inequality is satisfied*

$$\frac{w}{d} < \frac{x}{n}$$

If the two ratios are equal, then the planner will be indifferent between the two strategies.

.

To get a baseline comparison for what happens when windfall income is received, consider the situation where the planner is unaware of the illiquid savings option. The proportion of windfall income saved is dependent on the point in the savings cycle when it is received. The later in the savings cycle, or the earlier in the rest period, the less ego will be available in reserve and therefore the lower *s* will be. Additionally, all savings will be held in a liquid form, resulting in higher ego costs as we go further along in the savings cycle, therefore *s* will be inversely related to Y_t. Once the windfall has been received, ego will be depleted and so the rest cycle will begin.

The windfall has an equal probability of happening in each period, therefore the average *s* will be the sample mean of one savings/rest cycle. *s* in any given period is determined by equation 18. Assuming that the cost function has an inverse, we can solve for *s*

$$s = \omega^{-1}(R_t) - Y_t$$

Equation 21

And therefore average *s* will be

$$\bar{s} = \frac{\left(\sum_{t=1}^{d+w} \omega^{-1}(R_t) - Y_t\right)}{d + w}$$

But given that Y_t alternates between 0 and *y* during the rest period

$$\bar{s} = \frac{\left(\sum_{t=1}^{d} \omega^{-1}(R_t) - Y_t\right)}{d} + \frac{\left(\sum_{t=1}^{w} \omega^{-1}(R_t) - \frac{y}{2}\right)}{w}$$

Equation 22

When the planner is using the illiquid saving option, reaction to the windfall depends on the strategy being used. Satisfying equation 20 will allow the planner to maintain almost full ego and Y_t will be close to 0 since all savings will be in the illiquid option, so $\bar{s}$ will be a significant proportion of S. If equation 20 is not satisfied, then the amount of ego available will depend on the point in the saving/rest cycle when the windfall is received and on proposition 2. If $\frac{w}{d} > \frac{x}{n}$ then the planner will not fully deplete ego, but will have lower average ego than if equation 20 were satisfied. So $\bar{s}$ will be less than it would be if equation 20 were satisfied, but greater than in the situation where the planner is not putting savings into the illiquid asset or if $\frac{w}{d} < \frac{x}{n}$. In this final case the ego patterns will be similar to those where the planner is not aware of the saving asset, but since Y_t will be either *y* or 0 in every period, *d* will be greater and average ego levels will be higher than in the situation where there is no saving asset.

The windfall still has the same probability of occurring in each period. We find the sample mean of the values of s in each of the situations to determine average saving out of the windfall in each case. Going back to equation 21, we see that since liquid savings in each period will be 0 or y, s is given by $s = \omega^{-1}(R_t)$ or $s = \omega^{-1}(R_t) - y$ in alternating terms, and so $\bar{s}$ will be higher for all situations where the planner is using the saving asset. When equation 20 is not satisfied and *S2* is optimal, we have:

$$\bar{s} = \frac{\left(\sum_{t=1}^{d+w} \omega^{-1}(R_t)\right)}{d+w} - \frac{\sum_{t=1}^{\frac{d}{2}} y}{d}$$

$$= \frac{\left(\sum_{t=1}^{d+w} \omega^{-1}(R_t)\right)}{d+w} - \frac{\left(\frac{d}{2} * y\right)}{d}$$

$$\Rightarrow \bar{s} = \frac{\left(\sum_{t=1}^{d+w} \omega^{-1}(R_t)\right)}{d+w} - \frac{y}{2}$$

Equation 23

This equation differs from equation 22 for two reasons. First, the reduction of the Y_t term to $\frac{y}{2}$ means that $\bar{s}$ will be higher for any given ego level during the saving cycle. Second, because the ego costs are lower, *d* will be longer with the saving asset available resulting in more periods with a higher ego level than if it were not. Now suppose *S1* is optimal. By similar reasoning we then have that:

$$\bar{s} = \frac{(\sum_{t=1}^{x+n} \omega^{-1}(R_t))}{x+n} - \frac{y}{2}$$

Equation 24

Because ego is never depleted during this saving/rest cycle, equation 24 is equivalent to equation 23 where all terms with R_t below a minimum level are dropped, therefore $\bar{s}$ will perforce be higher than in equation 23. Now consider the situation where equation 20 is satisfied. There are few enough terms so that we can work out the sample mean term by term.

$$\bar{s} = \frac{\omega^{-1}(\Gamma) + (\omega^{-1}(\Gamma) - y)}{2}$$

$$= \frac{2\omega^{-1}(\Gamma) - y}{2}$$

$$\Rightarrow \bar{s} = \omega^{-1}(\Gamma) - \frac{y}{2}$$

Equation 25

Which is equivalent to equation 24 with all terms where ego reserve is not full dropped. Thus $\bar{s}$ will be highest in the situation where the planner is using the illiquid savings asset and equation 20 is satisfied.

This model has two major implications: first, precommitment eases self-regulation. By eliminating the choice for greater immediate gratification, the influence of the doer is significantly diminished and so the ability to self-regulate increases. Of course the problem with overregulation is that, just as with an overly strong planner, a commitment mechanism which is too strong can lead to overly restricted behavior and an inability to react flexibly to changing circumstances. Second, we replicate Fudenberg and Levine's result that more of windfall income is saved when precommitment is available. This seems inconsistent with the previous statement, but the difference is between commitment and overcommitment. In our model commitment is not overly strong. The individual is only committing one resource and only for a single period, although the planner ends up saving a great deal more than that.

C. Initialization Costs and the Importance of Framing

It is not uncommon for people to exhibit a sort of behavioral inertia. The larger a prospective task, the more intimidating it seems, and so the more difficult it is to initiate, hence the advice given in many self-help books to break large tasks into a series of smaller tasks and then handle them in sequence. This application will emphasize the importance of framing.

The individual's lifespan in this application is finite ($n < \infty$). In each period the individual can initiate an income generating activity which, after an initial period of disutility, yields income y for a number of periods q determined by the individual at the time of initiation. The individual must undertake the labor action in each of the q periods in order to receive the income, as well as to initiate the activity. As before, savings accrue interest at rate $i > 1$, labor and consumption are both actions which take the entire period to perform, but the illiquid saving asset is not available to the individual. If the individual fails to take the labor action in any round before the activity is complete, then the activity will end prematurely and a new one must be initiated before any more income can be generated. The activity yields an initial disutility $L(qy)$ which follows the restriction

$$\left|u\left(\frac{y(1-i^{q-1})}{1-i}\right)\right| < |L(qy)| < \left|u\left(\frac{y(1-i^{q})}{1-i}\right)\right|$$

This restriction ensures that in order to receive the positive utility from the activity, the planner must have sufficient ego in reserve to initiate the activity and take the labor action for all q periods. Given this restriction, utility from completing the activity will be

$$u\left(\frac{y(1-i^{q})}{1-i}\right) + L(qy)$$

Keeping in mind that L is negative since it is a loss function.

Given these restrictions, the planner's problem is then to maximize utility given the ego constraints from income and the initial disutility from initiation. The maximum number of rounds the planner can work given initial endowment R_1 is similar to equation 14, but note that in the initial period, the disutility from the ego cost will be

$$\omega(0 - L) = \omega(-L)$$

But since L is disutility, there will be a positive cost to the initiation period as well. If we place this cost into equation 14, then the maximum number of periods a planner can maintain assuming initial ego endowment Γ is the solution to the equation:

$$\sum_{k=1}^{q_{max}} \omega_{t+k} = \Gamma + q_{max} * r$$

Equation 26

It may not be in the planner's best interest to maximize q as this will involve a rest period afterward in order to build up ego supplies, using up the finite number of periods when the individual is active. The planner may instead wish to adopt a strategy with shorter income generating periods, but also shorter, or no rest periods[8]. The rest period given a q period activity will be

$$w = \left\lceil \frac{\Gamma - \omega\big(-L(qy)\big) - \sum \omega_t}{r} \right\rceil$$

A complete cycle will take one period for initiation, q periods of labor, one period for consumption, and w periods of rest. Therefore in an n period lifetime, the individual can perform $j = \left\lfloor \frac{n}{q+w+2} \right\rfloor$ activities, with one additional shorter activity using the remaining periods if the floor function leaves a remainder. The total utility will be

$$U_p = j * \left[u\left(\frac{y(1-i^q)}{1-i}\right) + L(qy) \right]$$

To find the optimal activity length, we take the derivative of this function with respect to q and set this equal to 0, then optimal q, denoted q^*, will satisfy the following equation

$$\frac{\partial U_p}{\partial q} = \frac{\partial j}{\partial q} * \left[u\left(\frac{y(1-i^q)}{1-i}\right) + L(qy) \right] + j * \frac{\partial \left[u\left(\frac{y(1-i^q)}{1-i}\right) + L(qy) \right]}{\partial q} = 0$$

$$\text{Where } 0 \le j \le n$$

Equation 27

Note that there may be a corner solution due to the restrictions on j. We cannot determine much more about the solution without specifying the utility, cost, and/or loss functions, but

[8] The situation with no rest periods is only possible if we drop the assumption that $\omega(y) > r$

the significance of this application is primarily in the interpretations rather than in the specific results of equation 27.

Actions are once again presented in terms of income for labor, but the motivation for this application comes from the idea that it is more difficult to initiate a task than it is to continue it. Suppose that instead of interpreting y as income, we view it as progress on an activity which will not yield any benefit until it is complete. Now suppose that the individual wishes to complete an overarching objective that will require Q periods of labor to finish. It is entirely possible that $Q > q_{max}$, in which case the individual will not be able to complete the objective in a single activity, and so may not be willing to attempt it since zero utility is better than the negative utility that would be gleaned from failing to complete an activity. However, if the planner instead tries to undertake a number of activities such that the total number of labor periods totals to Q, then the individual has a much better chance of completing the objective.

This interpretation underscores the importance of framing. If the larger task is framed in terms of a number of smaller sub-tasks that have to be completed, then it will be much easier and more likely to be completed. The framing need not be in terms of the size of the task, it could also be in terms of the time frame in which it needs to be completed. Suppose the goal is set by a manager for an employee to complete. The manager may have more success setting a number of shorter term deadlines that partly accomplish the overall goal than by simply telling the employee what needs to be accomplished in the long term.

Extensions

We now consider possible extensions or additions to the model which are suggested by the literature, but which were not necessary for the applications above. Arguments in this section are more intuitive than in our exploration of the applications. We first discuss the incorporation of additional factors into the cost function, as suggested by the experiments of Shiv and Fedorikhin (1999), and Leith and Baumeister (1996). We then consider the implications if the planner does not know the exact specification of the cost function, and instead must estimate it based on past experience. Finally we consider the addition of a utility cost to self-control.

A. Incorporation of Other Factors into the Cost Function

A number of experiments have shown that self-control is not the only activity which draws on ego. Shiv and Fedorikhin (1999) demonstrated that cognitive load increases impulsiveness and susceptibility to immediate temptation. Leith and Baumeister (1996) showed that physical fatigue reduces emotional regulatory capacity. There are two ways that we could model these effects. They could simply represent draws on the ego reserve, or they could be presented as a change in the parameters of the cost function. The actual difference in effect is largely, though not entirely, one of magnitude.

This extension is perhaps the most relevant of the ones considered here because it models the widest variety of situations, from the difficulty of making decisions while physically exhausted to dietary decisions of students during intense study periods. If we use the latter as an example, suppose a student is trying to choose between a candy bar and an apple as a study snack. The apple is obviously healthier, but the candy bar has more of an emotional appeal and is, therefore, more appealing to the doer. If the planner were to make a decision then the student would eat the apple, but because studying is effortful and depletes ego, the student ends up choosing the candy.

In the savings and labor investment application, modeling these alternative factors as a flat additional cost would have the same effect as reducing r. Because the cost from the environmental factor would not be dependent on y, it would be constant for every period when the outside factor is relevant, and so net ego income each period would be the difference between r and the additional cost. Therefore the effect would be a decrease in d and an increase in w. If we instead view the environmental factors as new parameters in the cost function, then the most likely effect they will have is to increase the marginal cost, similar to an increase in ω. In this case then the saving period would be shortened, but there would be no increase in the subsequent recovery time.

Once the planner is allowed the illiquid saving option, the changes depend on which equation determines $\bar{s}$. If equations 23 or 24 are relevant and the outside cost is constant, then d or x will be lowered respectively and w will be increased. This may have the effect of *raising* the average since there will be fewer periods where the planner has the additional cost of keeping the doer from consuming all savings. However, this effect will likely be quite weak and probably overshadowed by the additional costs from the environmental obstacles and the fact that lowered ego income will likely mean a lower overall average. If costs are not constant, then there will be no lowering of income, which would push $\bar{s}$ up, but it is easily possible that the environmental costs will be higher. If $\bar{s}$ is determined by equation 25, which is less likely due to the increased costs but still possible, then there will be no change in the behavior of the individual, but $\bar{s}$ will still be lowered by the additional cost. The difference between a constant increase and increased marginal costs depends entirely on the specifications of the environmental cost.

In the initialization cost situation, the effects of environmental costs are ambiguous since we are remaining agnostic as to the specification of the utility, loss, and cost functions. When q^* is a corner solution, the constant cost situation will reduce q^* and increase *w*, with the effect on *j* being dependent on the change in $q^* + w$. A decrease in *j* will push utility downward, and vice versa. $u\left(\frac{y(1-i^{q+1})}{1-i}\right)$ will decrease but so will $L(qy)$, so the change to the utility per activity will depend on the relative size of the changes in the two factors. If the increase is in marginal cost, then q^* will decrease with no commensurate increase in *w*, increasing j. The effect on utility per activity will still be ambiguous. When q^* is not a corner solution, then it will be the solution to equation 27, however *w* will increase in both the marginal and constant cost situations due to the increased ego costs.

B. Imperfect Knowledge of the Cost Function.

Now suppose that the planner does not know the exact specification of the cost function and must instead estimate it from previous experiences. This is similar to the self-reputation model of Benabou and Tirole and Kahneman's idea of remembered versus experienced utility. Each failure to self-regulate will lead to an increase in the estimated cost of self-regulation. We remain agnostic as to the exact form of estimation used by the planner, but it seems reasonable to assume that it is unbiased, and consistent. We could also assume stochastic costs, in which case, the estimate of the pdf will asymptotically converge to the actual pdf.

The implications for the savings and labor investment application are relatively simple. The inequalities in Proposition 1 will be calculated using the estimated mean of the cost function. If the estimated mean differs from the actual value then the consequences depend on which error the planner made. The planner will only need one failed saving cycle to realize that alternation will grant more utility. However if the planner is incorrectly alternating, then the planner will not realize the mistake until the estimate has converged sufficiently to the actual value. If costs are stochastic, then the “correct” strategy will be determined by the expected value of the cost function (i.e. the planner will be maximizing *expected* utility). A single failed saving cycle will not necessarily indicate that the planner has made the incorrect decision. Additionally, because the costs are not constant, it will take more periods for the estimate of the mean to converge sufficiently for the planner to realize the mistake.

In the precommitment situation with non-stochastic costs not much is different. A mistaken estimate regarding equation 20 will be resolved relatively quickly; the relative desirability of S1 and S2 will take a little more time since determining the optimal shorter saving cycle will be more difficult. The value of $\bar{s}$ will be unchanged, but unknown to the planner. If costs are stochastic, the changes are slightly more noticeable. s will be determined by the current cost level as well as the point in the saving cycle when it is received. Additionally, even if the costs satisfy equation 20 in expectation, there will be some cycles where they do not, and therefore there will be some periods where the planner defaults to strategy S1.

Estimation of costs is the most interesting part of the initialization cost application. Because the planner has to estimate whether or not there is sufficient ego remaining to complete an activity before initiating it, and an underestimation of cost can result in a loss, the planner’s best strategy will be to slightly overestimate costs in order to prevent an estimation error from resulting in an expensive mistake. This caution will be even greater in the case where costs are stochastic. This may mean that q^* is forced into being a corner solution which would lower utility and shorten the activities that the planner chooses to take.

As a result, the earlier point about the superiority of a strategy emphasizing frequent activities of shorter length is even more relevant.

C. Adding a Utility Cost to Self-control.

While not suggested by any empirical evidence, a number of theories assume that there is a utility cost to self-regulation over and above the lowered instantaneous utility from not consuming the maximal amount possible in a given period. We show here that we can add a utility cost into the model with few or no qualitative changes in our results. Our method is to add an additional term to instantaneous utility

$$\gamma\left[u(c_d) - u\left(c_p\right)\right]$$

which is a loss function increasing in the difference between the utilities of the doer and planner's consumption decisions. This has the effect of further reducing utility when $u\left(c_p\right) \notin argmax_d\, u(c_t)$.

In the savings and labor investment application, c_p is 0 in every period when the planner is exercising self-control, and so the loss function in period t will be

$$\gamma[u(c_d) - u(0)] = \gamma[u(c_d)] = \gamma[u(Y_t)]$$

And so when we substitute this into Condition 1 of Proposition 1, the saving cycle is the superior strategy if

$$u\left(\frac{y(1-i^d)}{1-i}\right) - \sum_{k=0}^{d-1} \sum \gamma[u(Y_{t+k})] > u(y)\left(\frac{1}{1-\delta^2} - \frac{1}{1-\theta}\right) * [1 - (\delta^{d+w})]$$

In other words, the saving cycle is less likely to be the superior option, and the additional utility is, unsurprisingly, lower than it would be otherwise. The equivalence of the two conditions still holds since the left side of both inequalities will be the same and all transformations showing their equivalence happened on the right hand side.

The utility cost has little or no effect on the results in the precommitment application. Since the planner is saving all income rather than consuming at regular intervals, and because the utility cost does not result in additional ego cost, the planner's behavior will not change at all. There is one possible exception to this rule due to the reason for the planner's obsessive saving behavior. The planner is saving for a massive consumption period at some point infinitely in the future. Since the disutility is effectively an infinite series, it may outpace the putative consumption period if it is sufficiently large and saving would then be undesirable even with the illiquid saving option.

The planner would also have the option of saving a large amount and then "living off the interest" by taking out only the interest income from the illiquid saving asset each period and then consuming this interest in the following one. This strategy would not require any self-regulation, but there are several problems. The biggest difficulty is deciding when to stop working and start consuming. If an individual stops working in any period then, they are giving up increased income by working for another period and then living off of the increased interest payments. There may be a third strategy where the planner takes regular periods to consume interest income until sufficient utility has been gleaned to compensate for the utility costs of self-regulation and then continues with the saving behavior.

The effects on initialization cost behavior can be quite extreme. If the costs from self-control are too high, then it may be the case that the utility cost from self-control is greater than the net utility from the activity for all possible values of q in which case the individual will undertake no activities at all. In a less extreme case, the utility cost may place an upper limit on q below that imposed by q_{max}. This limitation will only result in a change of behavior if the new limit is lower than q^* as determined by equation 27. If behavior does change, then utility will be lower than the case where utility cost is not a factor. It is worth noting that the utility cost from initiation is different from the utility cost to self-control. The initiation cost serves two purposes: it ensures that net utility is not positive until the activity has been completed, and it creates an initiation ego cost whose magnitude is dependent on the length of the activity in question.

Summary and Conclusion

Self-regulation is used in a number of situations, from eating to business negotiation. Because of this importance, there is a surfeit of work in the literature modeling and testing the mechanisms by which people control their behavior. Among those who work in the field, there is general agreement that both neurological and experimental evidence as well as theoretical deliberation indicate that decisions, especially those regarding self-regulation, are controlled by more than one neurological system. Impulsive decisions activate emotional areas of the brain, whereas those requiring more thought and consideration activate both the emotional areas and areas often associated with cognition and forward planning. Most models of the interplay between these various systems take the form of principal-agent models, referred to as dual-self models, where the emotional, more impulsive systems are represented by the agent and the cognitive decision systems take the part of the principal. In this paper, we adopt the planner and doer terminology used by Thaler and Shefrin in their seminal work on dual-self modeling.

The majority of dual-self models assume that the limiting factor on self-control is either inherent to the self-control method, as in the contract method of Brocas and Carillo, or a result of a utility cost imposed by self-regulation like Fudenberg and Levine or Thaler and Shefrin. One feature common to all of these models is that they assume that the planner can exercise as much self-control as desired if there is a contract that can be designed to effect it or if the planner is willing to pay the utility cost. The problem with this assumption is that there is a great deal of experimental evidence showing that self-control is in fact a limited resource.

Baumeister's results that persistence exercises reduce the ability to exert self-control clearly show that the self-regulation resource dubbed "ego" can be depleted over time, and once depleted, a rest period is needed in order to replenish it. Shiv and Fedorikhin demonstrated that high cognitive load will also impede self-control, meaning that this resource is drawn upon by a number of activities, not just self-regulation. These results form some of the foundations of the strength model of self-control in psychology. Despite the strong empirical base for this model, there is been little interest from behavioral economists. We designed our model to fill this gap by forcing the self-control method to draw on a limited reserve of ego.

In the savings and investment application, we show that the benefits of self-regulation increase with the individual's ability to self-regulate. This explains not only the need for rest breaks, but also why workers tend to be more productive if they are allowed regular rest periods rather than being forced to work throughout the day. One of the marks of a good theory is that its predictions can be tested in an experimental or econometric setting. Not only does our theory provide predictions for what factors influence whether a constant low level of work is more desirable than engaging in periods of more intensive work followed by a rest period, but it also makes predictions about the effects of the two possible ways that self-regulatory capacity could increase and how to differentiate them in an experiment. If subjects' rest periods are held constant, an increase in maximum ego capacity will not result in a subsequent increase in capacity for self-regulation. If practice decreases ego costs, then self-control will increase in this situation. Perhaps most importantly, we explain some of the benefits of having an impulsive system, as complete control on the part of the planner could lead to over-saving and under-consumption.

The second application is closer to a standard economic model in that it examines saving rates in the face of windfall income. Our results are not surprising as it makes sense that precommitment would leave more ego available for the planner to deal with unexpected draws. There is room for more work on the theoretical side of this application. We did not explore the effect of precommitment on increasing ability to exercise self-control. In part because we did not go into detailed definition of the process of ability increase, but even more because we did not discuss the effect of neglect on ability at all. It is entirely possible that too much precommitment may mean that the planner does not get enough practice at self-regulation; it may be possible for the planner to raise $\bar{s}$ by using strategy S1 or S2 and exercising the self-control "muscle" more often.

A similar criticism could be leveled at the initialization cost application. Increasing self-regulatory capacity means that the planner can pursue more ambitious activities, but may also have longer waiting periods in order to do so. Because of this, there may be little or no change in q^* with increasing Γ unless it is a corner solution, but it is almost certain that q^* will increase with a decrease in costs. This application is not closely related to most economic models since framing is more of a behavioural concern, but that very same quality also makes it more relevant to behavioral studies since they are, as previously noted, often concerned with framing. Our treatment of this application is weak on the experimental end since we do not fully solve for $q^{*,}$ and so we cannot state comparative statics to be tested in the lab.

If we are to state a general weakness in our theory, it is its difficulty in generalizing it to standard economic situations. Our saving application does not allow for halfway measures, and is, thus, not useful as a model describing saving rates in a general context. The reactions to windfall income and framing of tasks are closer to standard economics, but even they include measures that are difficult to quantify outside of a controlled experiment or possibly in a neuroeconomic context. On the other hand, these weaknesses are common to most dual-self modeling applications; both Fudenberg and Levine and Brocas and Carrillo have the same problems, but they highlight the strengths of dual-self models as well. While dual-self models tend not to produce results to be tested in the laboratory, they tell the story of the underlying mechanics behind behavior in more detail than most models. Other theorists can then draw upon the results of dual-self techniques in order to construct more specialized models which make testable predictions. By drawing upon laboratory work in its construction, our model creates a framework for future models to incorporate ego costs and thus reflect the reality of the underlying decision process.

References

Baumeister, R. F.; Bratslavsky, E.; Muraven, M.; Tice, D. M.. "Ego Depletion: Is the Active Self a Limited Resource?" *Journal of Personality and Social Psychology* 74 (1998): 1252–1265.

Baumeister, R. F., Vohs, K. D., & Tice, D. M. "The Strength Model of Self-control." *Current Directions in Psychological Science*, 16 (2007): 396-403.

Bechara, Antoine. "Decision Making, Impulse Control and Loss of Willpower to Resist Drugs: a neurocognitive perspective." *Nature Neuroscience*, 8 (2005): 1458 – 1463.

Benabou, Roland and Tirole, Jean. "Willpower and Personal Rules." *Journal of Political Economy*, 112, 4 (2004): 848-886.

Benabou, Roland & Pycia, Marek, "Dynamic Inconsistency and Self-control: a planner-doer interpretation." *Economics Letters*, Elsevier, 77, 3 (2002): 419-424.

Bodner, Ronit and Drazen Prelec. "Self-signaling and Diagnostic Utility in Everyday Decision Making." (2001) http://stellar.mit.edu/S/course/14/sp08/14.137/courseMaterial/topics/topic2/readings/bodnerprelec/bodnerprelec.pdf

Brocas I. and Carrillo, J. "The Brain as a Hierarchical Organization." *American Economic Review*, 98 (2008): 1312-1346.

Elster, Jon. "The Nature and Scope of Rational-choice Explanations." in *Actions and Events: Perspectives on the philosophy of Donald Davidson*, ed. Ernest LePore and Brian P. McLaughlin (Oxford: Blackwell, 1985): 60-72

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40, 2 (2002): 351–401.

Gailliot, M. T., Baumeister, R. F., DeWall, C. N., Maner, J. K., Plant, E. A., Tice, D. M., Brewer, L. E., & Schmeichel, B. J. "Self-control Relies on Glucose as a Limited Energy Source: Willpower is more than a metaphor." *Journal of Personality and Social Psychology*, 92 (2007): 325-336.

Kahneman, Daniel. "The Riddle of Experience vs. Memory." TED Video, 20:07, Mar 2010, http://www.ted.com/talks/daniel_kahneman_the_riddle_of_experience_vs_memory.html

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 62:443-477.

Leith, Karen Pezza, and Baumeister , Roy F. "Why Do Bad Moods Increase Self-defeating Behavior? Emotion, risk tasking, and self-regulation." *Journal of Personality and Social Psychology*, 71, 6 (1996): 1250-1267.

McClure, Samuel M., David Laibson, George Loewenstein and Jonathan D. Cohen.

"Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science*, 306 (2004)

McIntosh, Donald. *The Foundations of Human Society.* (Chicago: University of Chicago Press, 1969).

Muraven, M.; Tice, D. M.; Baumeister, R. F. "Self-control as a Limited Resource: Regulatory depletion patterns." *Journal of Personality and Social Psychology* 74, (1998): 774–789.

Muraven, M., Shmueli, D., and Burkley, E. "Conserving Self-control Strength." *Journal of Personality and Social Psychology*, 91 (2006):524-537.

Netzer, Nick. "Evolution of Time Preferences and Attitudes Toward Risk." *American Economic Review*, 99, 3 (2009): 937–955.

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

Schelling, Thomas C. "Self-Command in Practice, in Policy, and in a Theory of Rational Choice." *The American Economic Review*, 74, 2. (1984): 1-11.

Shiv, Baba and Alexander Fedorikhin. "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making." *Journal of Consumer Research*, 26, 3 (1999): 278-292.

Thaler, Richard H. and Shefrin H. M. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy*, 89, no. 2: 392-406.

Appendix: Proofs

Proof of proposition 1

We need to find the condition where utility from the savings cycle is greater than utility from alternation, i.e. that:

$$\frac{u\left(\frac{y(1-i^d)}{1-i}\right)}{1-(\delta^{d+w})}+\frac{u(y)}{1-\left[\frac{\delta^{d+2}-\delta^{d+w+2}}{1-\delta^2}\right]}>\frac{1}{1-\delta^2}u(y)$$

If we let $\theta=\left[\frac{\delta^{d+2}-\delta^{d+w+2}}{1-\delta^2}\right]$ to simplify notation then

$$U_p=\frac{u\left(\frac{y(1-i^d)}{1-i}\right)}{1-(\delta^{d+w})}+\frac{u(y)}{1-\theta}$$

and so our inequality becomes

$$\frac{u\left(\frac{y(1-i^d)}{1-i}\right)}{1-(\delta^{d+w})}+\frac{u(y)}{1-\theta}>\frac{1}{1-\delta^2}u(y)$$

Moving all $u(y)$ to the right hand side and factoring.

$$\frac{u\left(\frac{y(1-i^d)}{1-i}\right)}{1-(\delta^{d+w})}>u(y)\left(\frac{1}{1-\delta^2}-\frac{1}{1-\theta}\right)$$

which gives us condition 1 if we move the denominator to the right hand side. Now if we note that

$$\theta=\frac{\delta^{d+2}-\delta^{d+w+2}}{1-\delta^2}=\frac{\delta^{d+2}(1-\delta^w)}{1-\delta^2}$$

Applying this to the $1-\theta$ term

$$1-\left[\frac{\delta^{d+2}-\delta^{d+w+2}}{1-\delta^2}\right]=\frac{1-\delta^2}{1-\delta^2}-\frac{\delta^{d+2}(1-\delta^w)}{1-\delta^2}$$

combining fractions and factoring out δ^2

$$= \frac{1-\delta^2-\delta^{d+2}(1-\delta^w)}{1-\delta^2} = \frac{1-\delta^2\left(1+\delta^d(1-\delta^w)\right)}{1-\delta^2}$$

Which means that

$$\frac{1}{1-\theta} = \frac{1-\delta^2}{1-\delta^2\left(1+\delta^d(1-\delta^w)\right)}$$

Substituting this into the r.h.s of condition 1

$$u\left(\frac{y(1-i^d)}{1-i}\right) > u(y)\left(\frac{1}{1-\delta^2} - \frac{1-\delta^2}{1-\delta^2\left(1+\delta^d(1-\delta^w)\right)}\right)[1-(\delta^{d+w})]$$

Giving the quotients in the middle term a common denominator and then combining

$$\left(\frac{1-\delta^2\left(1+\delta^d(1-\delta^w)\right)}{(1-\delta^2)\left[1-\delta^2\left(1+\delta^d(1-\delta^w)\right)\right]} - \frac{(1-\delta^2)^2}{(1-\delta^2)\left[1-\delta^2\left(1+\delta^d(1-\delta^w)\right)\right]}\right)$$

$$\left(\frac{\delta^2-\delta^4-\delta^{d+2}(1-\delta^w)}{(1-\delta^2)\left[1-\delta^2\left(1+\delta^d(1-\delta^w)\right)\right]}\right)$$

Factor out δ^2 and we get

$$= u(y)\left(\frac{\delta^2\left(1-\delta^2-\delta^d(1-\delta^w)\right)}{(1-\delta^2)\left[1-\delta^2\left(1+\delta^d(1-\delta^w)\right)\right]}\right)[1-(\delta^{d+w})]$$

Which gives us condition 2:

$$u\left(\frac{y(1-i^d)}{1-i}\right) > u(y)\left(\frac{\delta^2\left(1-\delta^2-\delta^d(1-\delta^w)\right)}{(1-\delta^2)\left[1-\delta^2\left(1+\delta^d(1-\delta^w)\right)\right]}\right)[1-(\delta^{d+w})]$$