Measuring Self-Control

John Ameriks, Andrew Caplin, John Leahy, and Tom Tyler*

February 2004

Abstract

How significant are individual differences in self-control? Do these differences impact wealth accumulation? From where do they derive? Our survey-based measure of self-control provides insights into all three questions:

1. There are individual differences in self-control not only of a quantitative but also of a qualitative nature. In our sample, standard self-control problems of over-consumption are no more prevalent than are problems of *under*-consumption.
2. Standard self-control problems *do* impede wealth accumulation, particularly in liquid form. Problems of under-consumption have the opposite effects.
3. Self-control is linked to "conscientiousness", a personality trait much studied by psychologists. There is a related link with financial planning.

*The Vanguard Group (Ameriks), Department of Economics, New York University (Caplin and Leahy), and Department of Psychology, New York University (Tyler). Caplin thanks the Center for Experimental Social Science at New York University and the C.V. Starr Center at NYU for financial support. We would like to thank Xiaohong Chen, Douglas Fore, Faruk Gul, Guido Imbens, David Laibson, Julio Rotemberg, Ariel Rubinstein, and Yaacov Trope for helpful discussions. We gratefully acknowledge financial support for our survey provided by the TIAA-CREF Institute. All opinions expressed herein are those of the authors alone, and not necessarily those of The Vanguard Group, TIAA-CREF or any of either organization's employees, trustees, or affiliates.

1 Introduction

That self-control problems may in theory impede wealth accumulation has been understood for almost 50 years (Strotz [1956], Laibson [1997]). Yet empirical confirmation has been lacking. In this paper we use survey techniques to verify the importance of the self-control-wealth link. At the same time, we shed new light both on the profound individual differences in self-control, and on the underlying determinants of these differences.

Our survey-based approach to measuring self-control has roots in the findings of Mischel and his collaborators on delay of gratification (Shoda, Mischel, and Peake [1988]). Yet our precise formulation exploits recent advances in self-control theory, in particular the model of Gul and Pesendorfer [2001]. As described in section 2, we use an allocation scenario to elicit the level of self-control in this model. To a first approximation, self-control is measured as the difference between the intertemporal allocation initially viewed as optimal and the allocation that would be chosen in practice. Three findings concerning this measure of self-control stand out:

1. The current view of self-control problems as involving the need to suppress the immediate urge to consume is inadequate. In our sample, "present-bias" (the urge to consume today more than would be ideal) is no more prevalent than is "future-bias" (a tendency to consume less today than would be ideal), as shown in section 3.

2. We identify a robust relationship between measured self-control and the level of net worth, detailed in section 4. Those who believe that they will consume at a faster than ideal rate in our allocation scenario are less wealthy than those with the opposite beliefs. Self-control problems have a particularly powerful impact on the level of liquid wealth, in accordance with theoretical predictions.

3. Individual differences in self-control relate to deep differences in overall personality structure. The most well-researched measures of individual differences are the "Big Five" personality factors. One such factor, conscientiousness, is seen by psychologists as strongly related to self-control. In confirmation, we show in section 5 that high levels of conscientiousness reduce the scale of both problems of over-consumption and problems of under-consumption. This suggests that the positive association between planning (an aspect of conscientiousness) and wealth

accumulation identified by Lusardi [1999] may be intermediated by self-control (see also Ameriks, Caplin, and Leahy [2003a]).

Our findings rely on strong identifying assumptions concerning the interpretation of our survey questions. In section 6 we explore possible explanations for our results when these assumptions are false. None of these stand out as more compelling than our self-control based explanations.

Our results suggest that the impact of self-control problems on wealth accumulation is more intricate than is currently believed. Exploring the savings behavior of the very wealthy, Carroll [2000] hypothesized that their behavior could be understood only if they derived utility directly from wealth. Similarly, our results suggest that much wealth may be in the hands of "under-consumers": households who are tempted to accumulate rather than to consume. On a related note, Krusell, Kuruscu, and Smith [2002] calibrate a variant of the Gul-Pesendorfer model of asset pricing. They find that the model better fits with various asset pricing facts, such as the risk-free rate puzzle, if the temptation is to save rather than to consume.

Going beyond the sphere of wealth accumulation, our results suggest the need for researchers to broaden their outlook on self-control problems. Existing theories in both economics and psychology treat these problems as involving the inability to delay gratification. Yet our finding of future-bias suggests that self-control problems may underlie behaviors that have hitherto been seen in a different light. Additional research is needed not only on the behavioral implications of future bias, but also on its origins.

2 Survey Methods and Self-Control

We begin by outlining the existing empirical literature on self-control. We then outline our survey methodology, and the survey instrument itself. We close by discussing the sense in which our survey methodology is novel, and the sense in which it is familiar.

2.1 Literature Review

Laibson, Repetto, and Tobacman [2003] use data on wealth accumulation, credit card borrowing, and consumption-income co-movement to estimate a "representative" utility parameter governing self-control. In most specifications, they reject the null of exponential discounting in favor of the hyperbolic

discounting model of Laibson [1997]. Other efforts to estimate the effects of self-control problems include Fang and Silverman [2002] and DeJong and Ripoll [2003].

The paper by DellaVigna and Paserman [2001] is one of the few to assess the impact of *individual differences* in self-control. They use cross-sectional variation in self-control to predict variation in behavior.[1] They show that various indicators of self-control problems, such as smoking and contraceptive use, are negatively related to job search effort and exit rates from unemployment.

While Della Vigna and Paserman use behaviors to gauge individual differences in self-control, economists have also used questionnaire techniques. The traditional questions used to measure self-control involve asking how much money an individual would require at various future dates in place of a fixed immediate reward (Thaler [1981], Chapman [1996]). The implied discount rates are generally found to be far higher over short horizons, which is taken as evidence that the respondent has a self-control problem caused by over-weighting of the present. The extent of the self-control problem is assumed to be reflected in the extent of this over-weighting.

The "time versus money" approach to measuring self-control has various drawbacks, as pointed out by Frederick, Loewenstein, and O'Donoghue [2002]. One problem is that the answer may not be allocatively relevant. The identification of acceptance of the reward with a corresponding burst of consumption is tenuous: just because I accept an additional $100 today does not mean that I will immediately spend it on a fancy meal. There are many other problems: an expressed preference for the present may arise because immediate receipt of a reward is certain and convenient, while future receipt is uncertain and inconvenient; the implied self-control problem appears to be less significant when non-trivial rewards are at stake; rational individuals, even those with self-control problems, should use market interest rates to discount monetary amounts (Fuchs [1982]); and the implied level of self-control varies widely from experiment to experiment based on apparently irrelevant changes in context (Rubinstein [2000]). Rather than looking to clean up this procedure, we adopt an entirely different approach to measuring self-control inspired by the findings of the social psychologist, Walter Mischel.

[1] See also Passerman [2003].

2.2 Methodology

During the period 1968-1974, Mischel and his collaborators ran an experiment offering cookies to 550 four year old children of students and professors at Stanford University. Each child was asked whether he or she would prefer to have one cookie or two: the universal preference was for two. The attending psychologist then left the room, and the child was told that in order to get both cookies, he or she would have to wait until the psychologist returned. The child also had the option of calling the psychologist back into the room prematurely, and receiving only one cookie rather than two. After 20 minutes, the psychologist re-entered the room and appropriately rewarded any child who had been sufficiently patient.

Mischel conducted follow-up studies of parents in 1981-82 and in 1984, asking them to complete a personality survey for their child (the Adolescent Coping Questionnaire, ACQ), and to provide SAT scores. While the samples were small (ACQ information for 134, SAT scores for 94), the findings were striking. A strong relationship was found between the extent of the time delay in the initial experiment and the following two personality measures (Shoda et al. [1988]):

- How likely is your child to exhibit self-control in frustrating situations?
- How likely is your child to yield to temptation?

In both of the above cases, the coefficient in a univariate regression was significant at the 0.1% level. Just as significant was the relationship found between ability to delay gratification at age 4 and the Quantitative SAT score.

Mischel's results show that the ability to resist temptation is a very long lasting personality trait with significant cross domain validity. In addition, this ability correlates with other important life outcomes, in particular academic achievement.[2] The study also suggests the external validity of survey questions, such as those from the ACQ, based on intuitions concerning temptation and self-control. Our approach to measuring self-control draws

[2]Even social psychologists, famously reluctant to acknowledge that character structure displays any cross-domain regularities, have been forced to conclude that self-control differences may underlie differences in behavior and achievement over long periods of time. Ironically, the case against character variables as explanators of behavior is most strongly associated with Mischel himself (Mischel [1968]).

on these insights. Following Mischel, we hypothesize that self-control characteristics apply across many distinct domains, and that over time individuals become aware of their own level of self-control. To gain access to this knowledge, we present individuals with a hypothetical scenario involving possible temptation, and ask them to reflect on their ability to resist. To increase precision, we use the temptation and self-control model of Gul and Pesendorfer [2001] as a guide in framing our questions.

2.3 Temptation and Self-Control

Gul and Pesendorfer show that a specific assumption on preferences over choice sets, set betweenness, when added to other standard assumptions, gives rise to a new and fascinating class of utility functions. With these assumptions, observed behavior maximizes the sum of a standard direct utility function, and an additional "temptation"-based utility.

Consider a standard two period problem of allocating a fixed physical supply of a perfectly storable consumption good, W, across two periods. A Gul-Pesendorfer consumer maximizes the sum of a classical utility function and a second "temptation" function, $T(c_1, c_2)$. The utility associated with a particular set of feasible consumption choices, $A \subset \Gamma(W) \equiv \left\{(c_1, c_2) \in R^2_+ | c_1 + c_2 \leq W\right\}$, is:

$$V(A) = \max_{(c_1,c_2)\in A}[U(c_1, c_2) + T(c_1, c_2)] - \max_{(c_1,c_2)\in A}[T(c_1, c_2)],$$

$U, T : R^2_+ \rightarrow R$. Actual choices are made as a compromise between the "standard" utility function and the temptation function: the agent may be willing to move away from the otherwise ideal choice in order to reduce the disutility associated with rejecting the most tempting option.

A simple special case serves to clarify the workings of the model. Let both functions be logarithmic,

$$\begin{aligned} U(c_1, c_2) &= i \ln c_1 + (1 - i) \ln c_2; \\ T(c_1, c_2) &= \lambda[\tau \ln c_1 + (1 - \tau) \ln c_2]; \end{aligned}$$

with $0 < i, \tau < 1$, and $\lambda \geq 0$. In this case, the consumption profile most preferred by the individual as a singleton choice set involves consuming proportion i of the resource in the first period. On the other hand, with a larger choice set there is a temptation to consume a higher proportion τ in the first

period. With $A = \Gamma(W)$, the actual choice is a compromise between these two functions, giving weight $\frac{\lambda}{1+\lambda}$ to the temptation as opposed to the ideal choice. The actual proportion of wealth consumed in period 1 is therefore:

$$a = [\frac{1}{1+\lambda}]i + [\frac{\lambda}{1+\lambda}]\tau.$$

Our specific interest is in the level of self-control. This can be identified as the difference between the actual and the ideal proportion of wealth consumed in period 1:

$$a - i = (\frac{\lambda}{1+\lambda})(\tau - i). \quad (1)$$

Our goal is to design a question to measure this self-control parameter.

2.4 The Hypothetical Choice Problem

The Gul-Pesendorfer model maps well to psychological intuition concerning self-control. The parameter i measures the ideal split between current and future consumption, in the sense that it is the split the agent would most prefer with complete commitment. The parameter τ measures the most tempting allocation, in the sense that any deviation from it results in a utility penalty for "resisting temptation to move away from the ideal". The parameter λ characterizes the relative weight of the temptation in actual decisions, with lower values corresponding to a greater ability to resist temptation. Hence we sought to measure precisely these allocations in a hypothetical problem of resource allocation.

In designing our allocation problem, one key goal was to avoid ambiguities concerning the timing of consumption. With this goal in mind, we presented respondents with a scenario in which they had won a prize that they could use at any time in the next two years, but which would become valueless thereafter. We wanted the prize to be attractive, yet too expensive for most agents to pay for out of their own resources (to remove simple substitution into the general lifetime pattern of consumption). At the same time, we did not want the prize to be a completely indivisible once in a lifetime experience, since this would reduce the information content of our allocation question. Here is the precise scenario we used, including the words used to act as a bridge to the questions themselves:

- Suppose that you win 10 certificates, each of which can be used (once) to receive a "dream restaurant night." On each such night, you and a companion will get the best table and an unlimited budget for food and drink at a restaurant of your choosing. There will be no cost to you: all payments including gratuities come as part of the prize. The certificates are available for immediate use, starting tonight, and there is an absolute guarantee that they will be honored by any restaurant you select if they are used within a two year window. However if they are not used up within this two year period, any that remain are valueless.

The questions below concern how many of the certificates you would ideally like to use in each year, how tempted you would be to depart from this ideal, and what you expect you would do in practice:

- 3a. From your current perspective, how many of the ten certificates would you ideally like to use in year 1 as opposed to year 2?

- 3b. Some people might be tempted to depart from their ideal allocation in (a). Which of the following best describes you: (please mark only one)

    1. I would be strongly tempted to keep more certificates for use in the second year than would be ideal.
    2. I would be somewhat tempted to keep more certificates for use in the second year than would be ideal.
    3. I would have no temptation in either direction (skip to 3d)
    4. I would be somewhat tempted to use more certificates in the first year than would be ideal.
    5. I would be strongly tempted to use more certificates in the first year than would be ideal.

- 3c. If you were to give in to your temptation, how many certificates do you think you would use in year 1 as opposed to year 2?

- 3d. Based on your most accurate forecast of how you think you would actually behave, how many of the nights would you end up using in year 1 as opposed to year 2?

Our statistical analysis of the answers to these questions is based on the following three identifying assumptions.

- Identifying Assumptions
  - *A1*: We assume that the Gul and Pesendorfer model is valid.
  - *A2*: We assume that our question is answered in terms of the model, as a question concerning the value of the key parameters of the model in the allocation problem that is presented.
  - *A3*: We assume that the self-control parameter translates perfectly from our two period hypothetical choice problem to the more general problem of wealth accumulation.

Given these assumptions, our fundamental interest is in measured self-control, $a - i$, as defined in equation (1) above. Our identifying assumptions assert that this measure is identical in our free dinner scenario and in the general problem of wealth accumulation. Modulo the corner constraints discussed below, the gap between expected and ideal consumption corresponds precisely to the self-control problem in the Gul-Pesendorfer theory.

Note that while our question derives explicitly from the model of Gul and Pesendorfer, our measure of self-control may closely approximate true self-control even if alternative models of self-control apply. Consider for example the hyperbolic discounting model of Laibson [1997]. In this model, changing tastes give rise to a time inconsistency problem. The obvious interpretation of our measure of self-control in this model is that expected consumption is the solution to the game between the various temporal selves, while ideal consumption is the plan that maximizes utility from the present perspective.[3] In the planner-doer model of Thaler and Shefrin [1981], the doer values only current consumption, while the planner maximizes the present value of the utility of each doer. At cost, the planner can alter the doer's preferences to produce an interior optimum at a cost. In this model, ideal consumption would be the plan that maximized the planner's utility, while expected consumption would be the result of the interaction between planner and doer.[4]

[3] One difference between the models is that there is no obvious counterpart to temptation in the standard formulation of the hyperbolic model; the agent either commits to the ideal or adjusts current behavior in light of future choices. In addition, it is somewhat easier to rationalize the finding of under-consumption in terms of the Gul-Pesendorfer framework.

[4] It should not be surprising that we can reinterpret our question in terms of the planner-doer framework, since Benabou and Pycia [2003] show that the Gul-Pesendorfer model itself can be so interpreted.

Finally, Benhabib and Bisin [2002] and Bernheim and Rangel [2001] conceptualize self-control problems as involving conflict between the automatic and controlled pathways in the brain. The controlled pathways represent reasoned goal pursuit, while the automatic pathways represent programmed responses that reflect the influence of evolution or classical conditioning. In these formulations, ideal consumption would correspond to the reasoned optimum, while expected consumption would reflect also the influence of the automatic processes.

2.5 Methodology Revisited

In many ways our empirical methodology is standard, based as it is on strong but somewhat obvious identifying assumptions. Like us, Barsky, Juster, Kimball, and Shapiro [1997] use a theoretically-inspired approach to measure preference parameters such as the discount rate in the context of a hypothetical allocation problem. Yet there is one critical difference between our methodology and theirs. The earlier questions on preference parameters refer only to hypothetical choice experiments. In contrast our question on self-control refers not only to choices, but also to ideals and to temptation, two concepts that are not directly connected to a specific choice experiment.

While our effort to measure self-control using non-behavioral questions may be controversial, the work of Mischel is a powerful precedent. Apparently, intuitively derived measures of temptation and self-control may provide valuable insights into actual behavior. Of course, ambiguities remain. Temptation and ideals may mean different things to different people, making the answers to our question in practice somewhat hard to interpret. In formal terms, we get around this simply by relying on a well-specified theory, and making the identifying assumption that the question is interpreted in line with that theory. Yet our identifying assumption concerning the interpretation of a verbal question is particularly strong. This makes it incumbent upon us to consider seriously alternative interpretations of our empirical results given various alternative interpretations of the questions we ask. We address this issue in section 6 below, where we discuss the issues that may arise if our identifying assumptions are false.

Is it possible to ask questions on self-control that fit entirely within the classical choice theoretic framework? In principle, the answer is yes. After all, one of the beautiful features of the Gul-Pesendorfer model is that it is based entirely on choices that are at least potentially observable. This allows

us to pursue an alternative approach to measuring self-control based only on hypothetical choices. These choices concern the use of commitment devices. The precise questions and their answers are summarized in detail in section 4. When we use these questions as proxies for self-control, the results are in many ways similar to those using our primary temptation-based measure. However the evidence suggests that responses to these questions hinge as much on the psychology of commitment as they do on the psychology of self-control. Our temptation-based question avoids this form of contamination.

3 The Nature and Extent of Self-Control Problems

Our question on self-control was included in a new survey sent in February 2003 to a sample of TIAA-CREF participants. All of the 2,406 individuals who received the survey had responded to two previous surveys: the Survey of Participant Finances (henceforth SPF), fielded in January 2000; and the Survey of Financial Attitudes and Behavior (henceforth FAB), fielded in January 2001. Combining our three surveys, we have very rich data on personality, behavior, preferences, demographics, wealth, and income.

The response rate to our third survey was on the order of 68%, with some 1,632 providing responses. We removed 87 respondents from the sample who failed to answer both the questions on actual and ideal consumption. We also removed respondents for whom the answers to the question on self-control were clearly meaningless. In particular, we asked respondents to place a cash value on the free dinner prize, and removed from the sample the 25 for whom the prize had no value. Table 1 presents key demographic statistics for the remaining 1,520. The category totals in Table 1 are typically smaller than 1,520 due to non-response to individual questions.

Replying to the third survey has not fundamentally altered the demographic structure of the sample, although response rates were higher among older households. As before, respondents stand out in terms of their educational achievements, with roughly 1 in 3 being teaching faculty. In Ameriks, Caplin, and Leahy [2002] we compared financial characteristics of respondents to the first two surveys with those of working households in the 1998 Survey of Consumer Finances (SCF). Net worth is some 2.5–3 times higher in our sample, while debt levels are generally lower. Not only are the de-

mographic and economic profiles of respondents different from those of the general population, so too are their behavioral and psychological profiles. In particular, the sample is increasingly self-selected on the basis of interest in responding to intricate survey questions. For example, it is intuitively reasonable to expect survey respondents to be more conscientious than are non-respondents. As we will see in section 5 below, this extra conscientiousness may have significant implications for wealth holdings.[5]

3.1 Ideals and Expectations

Table 2 presents the distribution of answers to the questions concerning the ideal and expected allocation of resources. Some 60% of respondents indicated that their ideal allocation involved an equal split between the two periods. Among those who gave other answers, the overwhelming tendency was to wish to consume more in the first year, with eight times as many selecting answers of 6 and above than answers of 4 and below. The contrast at the extremes is especially striking. More than 15% of respondents stated a wished to consume all of their meals in the first year, with only a tiny fraction preferring to consume all in the second year.

The distribution of expected consumption is more dispersed than is that of ideal consumption. Less than 50% expect an equal split. There is also a greater tendency towards low consumption in the first year. Only three times as many select answers of 6 and above as opposed to answers of 4 and below.

3.2 Measured Self-Control

To a first approximation, the self-control parameter in the Gul and Pesendorfer model corresponds to the numerical difference between expected and ideal consumption, which we refer to as the EI gap. Table 3 in the appendix presents the joint distribution of expected and ideal consumption. Note that

[5] While our non-representativeness in economic and demographic terms clearly differentiates us from more standard surveys, not so the behavioral self selection. Respondents to surveys such as the HRS may be just as psychologically non-random as are respondents to our survey, since by definition they are the members of their demographic and economic cohort who were willing to answer the questions posed. As understanding of the role of behavioral variables advances, the methodology for achieving randomness in large national surveys will need to be amended.

while there are a few outliers, 95% of responses lie within two columns of the diagonal, implying that the EI gap is typically small. Note also that either the expected or the ideal consumption lies at a corner for about 17% of the observations. In technical terms, these corner observations imply that our measure of self-control is censored, and our statistical procedures are designed to handle this censoring as efficiently as possible.

A simple example clarifies the impact of corner constraints on our measure of the EI gap. Consider two individuals with identical self-control problems as regards wealth accumulation, yet different ideal levels of consumption. Individual A wishes ideally to consume 3 meals this year, in order to anticipate next year's meals with all the more pleasure (Loewenstein [1987]). Taking account of her self-control problem, she expects in fact to consume 7. Individual B is keener than is A to try new restaurants sooner rather than later, and he picks an ideal first year consumption level of 9. Given his self-control problem, he expects to consume all 10 in year 1. In this example, note that even though A and B have identical self-control problems, our survey fails to pick this up: A's EI gap is measured as 4, while B's is measured as 1. The corner constraint has censored our observation of B's self-control problem.

Given our identifying assumptions an observation may be censored only if either the expected or the ideal level of consumption is at a corner (10 or 0). In total, there are 267 such observations in our sample. Note that the most severe examples of censoring arise when both observations are at the same corner. In particular, there are some 123 individuals for whom expected and ideal consumption are both at the maximum value of 10. In these cases our measure of the EI gap has no information whatever on the nature of the underlying self-control problem. We will control for this censoring in the estimation.

Table 4 reports the distribution of the EI gap. The first column presents the distribution for all 1520 who answered the self-control questions. The third column presents the distribution for the 1253 for whom this measure is unaffected by the corner constraints.

Table 4 provides strong confirmation of our first finding. In our sample, more people expect to use less than their ideal number of certificates in the first year than expect to use more than ideal number. Hence problems of under-consumption are at least as prevalent as are those of under-consumption. The third column of table 4 which excludes ambiguous observations gives the clearest picture: almost one respondent in five has a problem of under-consumption, while only one in eight has a standard prob-

lem of over-consumption.

4 Self-Control and Wealth Accumulation

In this section we look first at the impact of self-control problems on net worth. We then look at wealth held in more and less liquid forms. Finally, we discuss the use of our alternative measure of self-control that is based entirely on choice behavior.

4.1 Self-Control and Net Worth

The simplest empirical approach to measuring the impact of self-control on wealth accumulation would be to run a standard linear regression,

$$w = \alpha_0 + \alpha_1 sc + \alpha_2' \mathbf{x} + \epsilon, \tag{2}$$

where w is net worth, sc is self-control and $\mathbf{x}$ is a vector containing other economic and demographic variables often included in classical life-cycle regressions. Yet our measure of self-control, the EI gap, may differ from the true underlying self-control measure sc due to the censoring problem described above. Using the EI gap to proxy for sc would bias α_1 away from zero, since censoring reduces the absolute value of the EI gap relative to sc. While discarding the censored observations would not produce bias, it would be very inefficient. A large self-control problem of either type makes censoring more likely, rendering censored observations among our most informative. Hence it is important to find a more creative solution to the censoring problem.

Our solution to the censoring problem is adapted from the imputation literature.[6] Suppose that certain values of sc were missing at random. We would first estimate the conditional distribution $f(sc|\mathbf{x})$. We would then replace each missing value with a draw from this distribution, and estimate (2). We would repeat this procedure a number of times and take as our estimate of α_1 the average of the estimated $\hat{\alpha}_1$'s.

Our procedure differs only in that we have some additional information regarding the missing observations. We know that the right censored observations are greater than the EI gap and the left censored observations are

[6] See Little and Rubin (2002) for a discussion of imputation. Whereas there is an immense literature dealing with the censoring of dependent variables, we know of no other papers dealing with the censoring of independent variables.

less than the EI gap. We therefore first estimate $f(sc|\mathbf{x})$. We regress

$$\text{EI gap} = \beta_0 + \beta_1'\mathbf{x} + \nu$$

Here we use all of the data including the censored observations, and we take account of the censoring in the estimation. Next we replace the censored observations with draws from $f(sc|\mathbf{x}, sc \geq \text{EI gap})$ or $f(sc|\mathbf{x}, sc \leq \text{EI gap})$ depending on the direction of the censoring. We repeat this procedure 10 times and take as our estimate of α_1 the average of the estimated $\hat{\alpha}_1$'s. The confidence intervals take into account imputation uncertainty and are calculated using standard multiple imputation techniques (Little and Rubin [2002], p.86).

Table 5 summarizes the results of this analysis. Our sample contains only 374 households, since we lack complete net worth data for many households. We also remove a few outliers with gross financial assets in excess of $5 million and exclude annuitants because their retirement wealth is difficult to assess.

Table 5 identifies a clear impact of self-control on wealth accumulation. Note that we include also the answer to question 3a on the ideal level of consumption, and find it to have no explanatory power whatsoever. In quantitative terms, the equation suggests that the average over-consumer accumulates some 18% less than one with no self-control problem, while the average under-consumer accumulates some 27% more.[7]

The finding of a powerful impact of self-control on wealth accumulation is very robust. Introducing additional right hand side variables, such as preference parameters, information on parental gifts and bequests, and wealth shocks, has little impact on the key finding. Dropping outliers or dropping censored observations also has little effect on the size or significance of the effect of the EI gap. As an additional check for robustness, we run a simpler procedure based only on a qualitative indicator of the nature of the self-control problem. Individuals with a strictly positive EI gap are classified as "over-consumers," while those with a strictly negative gap are classified as "under-consumers." Agents for whom ideal and actual are either both at the upper bound of 10, or both at the lower bound of zero are impossible to classify and therefore dropped from the analysis. The results are remarkably similar to those in table 5. An over-consumption problem reduces net worth

[7] The average gap among over-consumers in the sample used in the regressions is 1.4 certificates; the average gap among under-consumers is -2.1 certificates.

by 17%, whereas an under-consumption problem raises net worth by 27%. The two variables are jointly significant at the 5.4% level. Imposing the constraint that the over-consumption and under-consumption effects are equal and opposite yields a coefficient of .22 and a p-value of .017.

4.2 Self-Control and the Composition of Wealth

Most theories of self-control suggest that the impact of self-control problems on wealth should vary according to the liquidity characteristics of the underlying assets and debts. In particular, it should be hard for those with self-control problems to accumulate financial assets outside their retirement account. Yet with respect to retirement assets, even the sign of the effect of self-control problems is hard to predict, since the wealth-reducing urge toward immediate consumption may be offset by the desire to commit resources to the future in an illiquid form.

Table 6 confirms that there does indeed appear to be a more significant impact of self-control problems on liquid than on illiquid assets. The liquid assets we analyze are non-retirement financial assets. The less liquid assets are retirement assets. We restrict the sample to the group aged 64 and under because the difference in liquidity between retirement and non-retirement assets is radically reduced when individuals reach the age of retirement. We include agents for whom we have data on these components of wealth, but lack data on total net worth. The resulting sample contains 362 households and is the same for both regressions.

As theory predicts, the impact of self-control problems on non-retirement financial assets is larger and more statistically significant than that on illiquid retirement assets. A one-standard-deviation increase in the EI Gap reduces non-retirement financial assets by roughly 31%, and retirement assets by 12%.[8] The latter effect is statistically insignificant.

4.3 A Choice-Based Measure

As already discussed, we pursued an alternative approach to measuring self-control that was based purely on (hypothetical) choices. To this end our earlier question on temptation and self-control were followed immediately by questions concerning the choice of commitment devices:

[8]The standard deviation of the measure in this sampe is 1.52 certificates.

- 3e. Suppose that you had the option to restrict some of the certificates for use only in the second year. Would you use this option?
  - Yes
  - No
- 3f. Suppose that you had the option to restrict some of the certificates for use only in the first year. Would you use this option?
  - Yes
  - No

If respondents answered yes to either question 3e or question 3f, they were asked in addition to specify precisely how many certificates they would so restrict.

The answers to these questions define a set of constrained allocations that the respondent finds desirable. Agents have self-control problems if their expected consumption lies outside this set. In this case our measure of the self-control problem is the signed distance between the expected choice and the constraint set. We call this the revealed preference (RP) gap. As before this measure allows for both types of self-control problem, and is zero in the absence of such a problem.

Table 7 shows the distribution of the RP gap. We have fewer observations than for the EI gap, due mainly to non-response. We also eliminate a small number of observations for which the constraint set is empty. Note also that only one individual is impacted by corner constraints.

Two differences between the RP gap and the EI gap stand out. First, there are fewer agents with self-control problems when measured by the RP gap. Just over 10% of agents have self-control problems as measured by the RP gap, as opposed to just over 30% according to the EI gap. Second, the RP gap sheds a different light on the relative importance of the two types of self-control problems. According to the RP gap, those with self-control problems are more likely to have problems of over-consumption than of under-consumption, whereas under-consumption was more prevalent according to the EI gap.

In spite of the differences between the EI gap and the RP gap, the two measures are highly correlated. The correlation is .24. Moreover, when we use the RP gap in place of the EI gap in our wealth regression, we again

find a negative and significant relationship between self-control and wealth accumulation. The coefficient on the RP gap is -.195, with a standard error of .092, and significant at the 4 percent level.[9] The relationship between wealth and the RP gap, however, is based on only 32 individuals with non-zero self-control problems. Given the small numbers, the magnitude and significance, but not the sign, of the estimated effects are very sensitive to changes in the sample. This was not the case with the EI gap.

While our revealed preference measure of self-control does correlate to a certain extent with wealth, it may be corrupted by considerations that are external to the model. Specifically, the psychology of flexibility and that of self-control are somewhat distinct, and they get inextricably linked in our revealed preference question. Some respondents may dislike flexibility, without this having any obvious connection to self-control problems. For example, the vast majority (about 80%) of the respondents who restrict certificates to one period also choose to restrict some to the other period, possibly to enhance planning for and anticipating the meals.[10] On the other side, even individuals who have self control problems, yet who also have a desire to retain flexibility, may report a zero RP gap. In support of this, those with self-control problems according to the EI gap *but not* according to the RP gap are much less likely to set budgets and to plan for vacations than are those with both types of self-control problem.

5 What Determines Self-Control?

One of the main reasons for measuring individual differences in self-control is to improve our understanding of how these differences arise. In this section we look to personality psychology for first insights on the determinants of self-control. Specifically we consider the impact on self-control of conscientiousness, one of the Big Five traits uncovered by personality psychologists. In section 5.1 we provide a brief introduction to the Big Five model and to conscientiousness in particular. In section 5.2 we introduce our survey measures of conscientiousness, and demonstrate the subtle yet intuitive connection between conscientiousness and self-control. In section 5.3 we present

[9] This regression included the same variables as the wealth regression with the EI gap, with the exception of ideal consumption. There were 340 observations.

[10] It may also be the case that respondents did not understand that by restricting some certificates to one period they did not have to restrict some to the other.

preliminary findings on two other influences on self-control, one demographic, and one related to behavior. The demographic variable that we discuss is age, while the behavioral variable is financial planning.

5.1 The Big Five

In a recent survey article in the Handbook of Personality, Oliver John and Sanjay Srivastava write of personality psychology:

> "After decades of research, the field is finally approaching consensus of a general taxonomy of personality traits, the "Big Five" personality dimensions." [John and Srivastava, 1999, p. 103]

The Big Five represent personality at a very broad level of abstraction, with each summarizing a large number of distinct, more specific personality characteristics. The five factors are typically given numerical and linguistic labels as follows: factor I, extraversion; factor II, agreeableness; factor III, conscientiousness; factor IV, neuroticism: factor V, openness.

Even a cursory look at the Big Five suggests that the characteristic most closely related to self-control is conscientiousness, since all of the others deal with entirely different aspects of the personality. It turns out that this connection is far deeper than one might expect simply from the dictionary definitions. In their influential schema, Costa and McCrae [1992] break down each of the Big Five into a number of different facets: one of the six facets of conscientious individuals is that they are not impulsive. Similarly, John and Srivastava [1999] provide a more expansive definition of factor III as follows:

> "Conscientiousness describes socially prescribed impulse control that facilitates task- and goal-directed behaviors, such as thinking before acting, delaying gratification, following norms and rules, and planning, organizing, and prioritizing tasks" (p.121).

Given the close link between conscientiousness and self-control, psychologists have studied the role of conscientiousness in various patterns of behavior associated with lack of self-control. In a study of adults, conscientiousness has emerged as the only general predictor of job performance. In a study of adolescents, John et al. [1994] found that low levels of conscientiousness predict juvenile delinquency and various other disorders, while high levels were associated with superior school performance. As John and Srivastava point

out, one of the long term goals of this research is to help design interventions that might "teach children low in conscientiousness relevant behaviors and skills (e.g. strategies for delaying gratification)" (p. 125). Finally, research by Friedman et al. [1995a, b] and by Roberts and Bogg [2003] suggests that conscientious individuals have superior health and longevity outcomes.

5.2 Conscientiousness and Self-Control

We included a series of questions in our survey to allow the link between conscientiousness and self-control to be studied. The particular questions we used were drawn directly from the questionnaires presented in Costa and Widiger [1994]. Respondents were asked to indicate on a simple six point scale the extent of their agreement or disagreement with the following statements.

- Q1g: Sometimes I am not as dependable or reliable as I should be.
- Q1h: I never seem able to get organized.
- Q1i: I often feel that I speak or act too quickly, without thinking about the consequences.
- Q1j: I am often late for appointments.

We begin with a negative finding. Our measures of conscientiousness have essentially no correlation with the raw EI gap. If we regress the level of the EI gap on the variables that we included in the wealth regression and include our measures of conscientiousness, none of the conscientious measures are remotely significant, either individually or as a group.[11] More broadly, the raw EI gap appears uncorrelated with any variables of economic, demographic, or psychological interest.

On reflection, this negative finding is entirely consistent with the personality theoretic view of the relationship between conscientiousness and self-control. The Big Five approach does not suggest that there should be a simple monotonic relationship between conscientiousness and the EI gap. Rather, it suggests that conscientiousness should shrink the EI gap toward zero and away from either extreme positive or extreme negative values. An individual

[11] This regression used all 976 observations for which complete data on all of the variables was available.

with a large positive EI gap has a standard problem of over-consumption. The natural prediction is that those who are highly conscientious should have less significant self-control problems of this type. Algebraically, this corresponds to a smaller EI gap. On the other hand, a large negative EI gap indicates a self-control problem of under-consumption. Again, a conscientious individual should have a less significant self-control problem of this sort. Algebraically, this corresponds now to a larger (less negative) EI gap.

When we explore the data from this viewpoint, we find that conscientious individuals do indeed appear to have lesser problems of self-control. For those who are highly conscientious, there is a lower divergence between actual and ideal consumption, regardless of sign. For example, the standard deviation of the EI gap is 1.79 for the 124 respondents who indicated agreement to both questions 1g and 1h, and only 1.05 for the 962 respondents who indicated disagreement on both questions. In essence, the conscientious have smaller problems of self-control in either direction than do those who are not conscientious.

In table 8, we present evidence that conscientiousness is a good predictor of whether or not an individual has a self-control problem. The table contains results of a probit regression where the variable to be predicted is an indicator of a self-control problem (a non-zero EI gap). The explanatory variables include the first two of our conscientiousness questions (Q1g and Q1h), concerning respectively dependability and organization (all four measures are significant when included individually; these are the two questions that retain significance when all four are added to the right-hand side together). In order to increase sample size, we add only age, gender, and the ideal level of consumption as explanatory variables, ending up with a sample of 1300. Both of our measures of conscientiousness are strong predictors of whether or not an individual has a self-control problem. The results are completely robust to alternative specifications.

This relationship between self-control and conscientiousness works separately for problems of over-consumption and for problems of under-consumption. If we limit our regression to the 1015 observations for which the EI gap is non-negative, then the coefficient on Q1g is .10 with a p-value of .015 and the coefficient on Q1h is .12 with a p-value of .005. If we limit our regression to the 1133 observations for which the EI gap is non-positive, the coefficient on Q1g is .05 with a p-value of .258 and the coefficient on Q1h is .16 with a p-value less than .001.

5.3 Other Influences on Self-Control

Table 8 contains an interesting finding concerning the impact of age on self-control problems. There is a profound reduction in the scale of these problems as individuals age. Again, this finding shows up only when one uses the indicator for existence of a self-control problem, not the level of self-control. Older individuals experience fewer self-control problems either of over-consumption, or of under-consumption, than do their younger counterparts. This finding is certainly consistent with the psychological literature, in which it is a common-place that temptation falls with age.

The strong links that we have uncovered between conscientiousness and self-control shed new light on the findings of Ameriks, Caplin, and Leahy [2003a] concerning the relationship between wealth accumulation and the "propensity to plan". As one might expect, there is a strong relationship between the propensity to plan and conscientiousness: the extent to which the agent enjoys planning for vacations is highly correlated in our sample with both measures of conscientiousness, as well as with the indicator of a non-zero EI gap. This suggests that an increase in the propensity to plan will increase wealth accumulation *only* for individuals with the standard problem of over-consumption. For those with an under-consumption problem, increases in the propensity to plan should have the opposite effect of lowering wealth accumulation. The net effect of the propensity to plan on wealth accumulation may depend on the mixture of self-control problems that are present in the general population.

6 Alternative Explanations

We examine possible explanations for our results when our identifying assumptions fail. Most of the issues concern the interpretation of the answers to our survey question. Are there plausible interpretations of the answers to our questions that would rationalize our findings even if self-control has absolutely no real impact on wealth accumulation?

6.1 Social Desirability and Rationalization

It is possible that individuals who have high wealth for reasons that have nothing to do with self-control answered our questions *as if* they had problems of under-consumption. One possible explanation for this behavior might

be social desirability. Looking at our questions, they might have understood that we were looking for just such a correlation, and their desire to conform might have done the rest. A second possible explanation for such behavior would be rationalization: looking at their own high wealth, they may have concluded that they cannot be the types who consume more than they would ideally like to.

While it is not possible for us to provide strong evidence against this alternative hypothesis, there are reasons to doubt its power.

- Our data provide no evidence to support the idea that high wealth causes individuals to vary their reported level of self-control. If such an effect were present, we would expect exogenous shocks to wealth uncorrelated with true self-control to shift reported self control. Our latest survey contains explicit measures of two such exogenous shocks to wealth: an indicator for the past receipt of an unexpected gift or bequests, and an indicator for major unexpected expenses. These two variables are of the expected sign, and are jointly significant at the 2% level in our basic wealth regression. Yet their impact on the measured level of self-control turns out to be statistically and economically negligible.

- If measured self-control is simply an artefact of actual wealth seen through a filter of social desirability or rationalization, then what explains the relationship between self-control and conscientiousness? The fact that this relatively subtle relationship shows up so clearly in our data makes it highly plausible that our measure of self-control is correlated with true self-control. A proponent of the social desirability explanation might argue that while both correlations are present, the explanation for the measured self-control-wealth correlation has nothing to do with the measured self-control-real self-control correlation. This is a fine line to tread.

- If social desirability and self-justification were the only forces that explained the correlation between wealth and measured self-control, then one would expect survey measures of other preference parameters such as the discount factor and the precautionary motive to correlate highly with actual wealth accumulation. Yet the history of questions on these subjects is in large part a history of irrelevance. For example, we have in a previous survey asked a detailed question concerning the discount

factor, in which the desired answer was far more directly visible, and rationalization far closer to hand. Yet there was absolutely no correlation between the measured discount factor and wealth accumulation. In the current study, one might have expected social desirability and self justification to produce a correlation between actual wealth and impatience, as measured by the expected amount of consumption in the first year. Yet this relationship is extremely weak, and is reduced to nothing in the presence of measured self-control.

6.2 The Ideal World

There is considerable scope for misinterpretation of question 3(a) on ideal consumption. In particular, those who are currently more busy than they would like might ideally wish to consume more in the first year than they believe they will. Conversely, those with a surfeit of free time may prefer a world in which they were busier and went out less. In both cases, the ideal represents what the respondent would do in a better environment.

If the answers have an "ideal world" interpretation, our self-control problem of over-consumption may in part reflect a quite different problem. The interpretation of the correlation with wealth would be that those who were too busy at work or with the family to go out as much as they would like have accumulated unusually large amounts of wealth, while those who are not as busy as they would like have accumulated little wealth. While such a story can be told, the correlation with conscientiousness seems harder to explain. The interpretation would have to be that conscientious individuals for some reason undergo smaller fluctuations in either direction in how busy they are, or that even though they may be extremely busy, they nevertheless expect to go out just as they would have if they were less busy. Possible, but hardly compelling.

6.3 Absent-Mindedness

The answers to our questions on self-control may be impacted by absent-mindedness of a sort analyzed by Piccione and Rubinstein [1997]. Absent-minded individuals know that they are forgetful, and hence may expect not to use all of their certificates over the two year period. Ideally, consumers of this sort may wish to use all 10 certificates in two years: 5 in the first year and five in the second. Yet they may believe that they will forget about the

certificates until it is too late, and end up using only 8 of the certificates, 4 in each year. What our data would record in such cases would be self-control problems of under-consumption. The real issue would be a problem of control, which we would instead treat as a problem of self-control.

Even if absent-mindedness is at play in the answers to our questions on self-control, the connection between absent-mindedness and the answer to question 3 is likely to be subtle. In the example in the last paragraph, an absent-minded individual consumed less than the ideal number of certificates in the first year. Yet some absent-minded individuals may choose instead to consume *more* than the ideal number of certificates in the first year, in order to reduce possible waste caused by later forgetfulness. Another important caveat to this case is that the connection between wealth accumulation and absent-mindedness is little understood. Ameriks, Caplin, and Leahy [2003b] have analyzed the optimal pattern of consumption and savings for an absent-minded consumer. While the comparison is complex, the results suggest that those who are absent-minded generally have an incentive to consume more rather than less than those with perfect memories. Hence our finding that those with negative EI gaps are generally wealthy argues that at least some of this effect is explained by problems of self-control rather than problems of control.

7 Concluding Remarks

We have used survey techniques to generate new insights into the nature and implications of self-control problems. Clearly, self-control problems represent a fascinating link between psychological forces and economic behavior, and survey techniques have much to offer to our search for understanding of cause and consequence.

References

[1] Ameriks, John, Andrew Caplin, and John Leahy, 2002, "Retirement Consumption: Insights from a Survey," NBER Working Paper No. 8735.

[2] Ameriks, John, Andrew Caplin, and John Leahy, 2003(a), "Wealth Accumulation and the Propensity to Plan," *Quarterly Journal of Economics*.

[3] Ameriks, John, Andrew Caplin, and John Leahy, 2003(b), "The Absent-Minded Consumer," New York University working paper.

[4] Barsky, Robert, Miles Kimball, Thomas Juster, and Matthew Shapiro, 1997, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach Using the Health and Retirement Survey," *Quarterly Journal of Economics* CXII, 537-579.

[5] Benabou, Roland, and Marek Pycia, "Dynamic Inconsistency and Self-Control: A Planner-Doer Interpretation," *Economic Letters*,

[6] Benhabib, Jess, and Alberto Bisin, 2002, "Self-Control and Consumption-Savings Decisions: Cognitive Perspectives", Working Paper, New York University.

[7] Bernheim, Douglas and Antonio Rangel, 2001, "Addiction, Cognition, and the Visceral Brain", Working Paper, Stanford University.

[8] Carroll, Christopher, 2000, "Why Do the Rich Save So Much?", in Joel Slemrod, ed., *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, Cambridge:Harvard University Press.

[9] Chapman, G., 1996 "Temporal Discounting and Utility for Health and Money," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 22, 771-791.

[10] Costa, P.T., and R.R. McCrae, 1992, *NEO PI-R Professional Manual*, Odessa, Florida, Psychological Assessment Resources.

[11] Costa, Paul T., and Thomas A. Widiger, (Eds.), 1994, *Personality disorders and the five-factor model of personality*, Washington, D.C.: American Psychological Association.

[12] DeJong, David, and Marla Ripoll, 2003, "Self-control Preferences and the Volatility of Stock Prices," Working paper, University of Pittsburgh.

[13] DellaVigna, Stefano, and Daniele Paserman, 2001, "Job Search and Impatience", Working Paper, University of California at Berkeley.

[14] Digman, J.M., 1990, "Personality structure: the emergence of the five factors model", *Annual Review of Psychology*, 41, 417-440.

[15] Fang, Hanming, and Dan Silverman, 2002, "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY".

[16] Frederick, S., G. Loewenstein, and T. O'Donoghue, 2002 "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, vol. 40, 351-401.

[17] Friedman, Howard, Joan Tucker, Joseph Schwartz, Leslie Martin, Carol Tomlinson Keasey, Deborah Wingard, and Michael Criqui, 1995, "Childhood Conscientiousness and Longevity: Health Behaviors and Cause of the Death", *Journal of Personality and Social Psychology*, 68, 696-703.

[18] Friedman, Howard, Joan Tucker, Joseph Schwartz, Carol Tomlinson Keasey, Leslie Martin, Deborah Wingard, and Michael Criqui, 1995, "Psychosocial and Behavioral Predictors of Longevity", *American Psychologist*, 50, 69-78.

[19] Fuchs, V., 1982, "Time Preference and Health: An Exploratory Study", in *Economic Aspects of Healt*h, edited by V. Fuchs, Chicago, University of Chicago Press.

[20] Gul, Faruk, and Wolfgang Pesendorfer, 2001, "Temptation and self-control", *Econometrica* LXIX, 1403-1435.

[21] John, O.P., A. Caspi, R.W. Robins, T.E. Moffitt, and M. Stouthamer-Loeber, 1994, "The "Little Five": Exploring the Nomological Network of the Five-Factor Model of Personality in Adolescent Boys", Child Development, 65, 160-178.

[22] John, Oliver P., and Sanjay Srivastava, 1999, "The Big Five Trait Taxonomy," chapter 4 of L.A. Pervin and O. John (eds.), *Handbook of Personality,* New York, Guilford Press.

[23] Krusell, Per, Burhanettin Kuruscu, and Anthony Smith, 2002, "Time Orientation and Asset Prices", *Journal of Monetary Economics* 49, 107-135.

[24] Laibson, David, 1997, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* CXII, 443-477.

[25] Laibson, David, Andrea Repetto, and Jeremy Tobacman, 2003, "Wealth Accumulation, Credit Card Borrowing, and Consumption-Income Comovement," Working paper, Harvard University.

[26] Little, Roderick, and Donald Rubin, 2002, *The Statistical Analysis of Missing Data*, Wiley Interscience.

[27] Loewenstein, G., 1987, "Anticipation and the Valuation of Delayed Consumption," *Economic Journal*, 97, 666-684.

[28] Lusardi, Annamaria, 1999, "Information, Expectations, and Savings," ch. 3 in *Behavioral Dimensions of Retirement Economics*, Herry Aaron, ed. (New York: Brookings Institution/Russell Sage Foundation), pp. 81-115.

[29] Mischel, W., 1968, Personality and Assessment, New York, Wiley.

[30] Paserman, Daniele, 2003, "Job Search and Hyperbolic Discounting: Estimation and Policy Evaluation", Working Paper, Hebrew University.

[31] Piccione, Michele, and Ariel Rubinstein, 1997 "On the Interpretation of Decision Problems with Imperfect Recall", *Games and Economic Behavior*, 20, 3-24.

[32] Roberts, B. W., & Bogg, T., 2003, "A 30-year longitudinal study of the relationships between conscientiousness-related traits, and the family structure and health-behavior factors that affect health," *Journal of Personality*, in press.

[33] Rubinstein, A., 2000, "Is it "Economics and Psychology"?", forthcoming in *European Economic Review*.

[34] Srivastava, Sanjay, Oliver John, Samuel Gosling, and Jeff Potter, 2003, "Development of Personality in Early and Middle Adulthood: Set Like

Plaster or Persistent Change?", *Journal of Personality and Social Psychology*, 84, 1041-1053.

[35] Shoda, Y., W. Mischel, and P. Peake, 1988, "Predicting Adolescent Cognitive and Self-Regulatory Competencies from Preschool Delay of Gratification: Identifying Diagnostic Conditions", *Developmental Psychology*, 26, 978-986.

[36] Strotz, Robert, 1956, "Myopia and Inconsistency in Dynamic Utility Maximization, *Review of Economic Studies*, 23, 165-180.

[37] Thaler, R., 1981, "Some Empirical Evidence on Dynamic Consistency", *Economic Letters*, 201-207

[38] Thaler, Richard, and H.M. Shefrin, 1981, "An Economic Theory of Self-Control", *Journal of Political Economy* 89, 392-406..

Table 1
Demographic Characteristics of
2003 Survey Respondents

Characteristic	(n)	(%)
Gender		
Female	658	43.7
Male	849	56.3
Marital Status		
Curr. married	1016	67.6
Prev. married	278	18.5
Never married	210	14.0
Education		
College or below	429	28.2
Masters or Prof.	533	35.1
Ph.D.	558	36.7
Occupation		
Teaching faculty	538	40.8
Mgmt., Sen. Admn.	242	18.3
Other Tech./Prof.	276	20.9
Other	264	20.0
Age		
Below 35	79	5.3
35-44	156	10.4
45-54	285	19.0
55-64	368	24.5
65-74	438	29.1
75+	177	11.8
Number of children		
0	1167	77.4
1	131	8.7
2	162	10.8
3+	47	3.1

Source: Authors' tabulations of 2003 FAB survey data.

Table 2
Ideal and Expected Allocation of Certificates to First Year

	Ideal		Expected	
Certificates	(n)	(%)	(n)	(%)
0	3	0.2	4	0.3
1	10	0.7	14	0.9
2	15	1.0	26	1.7
3	13	0.9	62	4.1
4	31	2.0	114	7.5
5	907	59.7	713	46.9
6	192	12.6	243	16.0
7	58	3.8	77	5.1
8	43	2.8	55	3.6
9	0	0.0	6	0.4
10	248	16.3	206	13.6
Total	1520	100.0	1520	100.0

Source: Authors' tabulations of 2003 FAB survey data.

Table 3
Cross-tabulation of Ideal and Expected
Allocation of Certificates to First Year

	Expected											
Ideal	0	1	2	3	4	5	6	7	8	9	10	Total
0	3	0	0	0	0	0	0	0	0	0	0	3
1	0	10	0	0	0	0	0	0	0	0	0	10
2	1	0	11	1	1	0	1	0	0	0	0	15
3	0	0	3	10	0	0	0	0	0	0	0	13
4	0	0	2	6	20	2	1	0	0	0	0	31
5	0	1	8	38	76	643	99	28	4	1	9	907
6	0	0	0	5	16	40	118	6	6	1	0	192
7	0	0	0	1	0	10	9	31	4	0	3	58
8	0	0	1	0	0	4	8	4	22	1	3	43
9	0	0	0	0	0	0	0	0	0	0	0	0
10	0	3	1	1	1	14	7	8	19	3	191	248
Total	4	15	28	65	118	718	249	84	63	15	216	1,520

Source: Authors' tabulations of 2003 FAB survey data.

Table 4
The Expected-Minus-Ideal Gap

	All Observations		Uncensored Observations	
E-I	(n)	(%)	(n)	(%)
5	9	0.6	0	0.0
4	2	0.1	2	0.2
3	8	0.5	5	0.4
2	39	2.6	36	2.9
1	113	7.4	113	9.0
0	1,059	69.7	865	69.0
-1	141	9.3	138	11.0
-2	94	6.2	74	5.9
-3	25	1.6	17	1.4
-4	9	0.6	2	0.2
-5	14	0.9	0	0.0
-6	2	0.1	1	0.1
-7	1	0.1	0	0.0
-8	1	0.1	0	0.0
-9	3	0.2	0	0.0
Total	1,520	100.0	1,253	0.0

Source: Authors' tabulations of 2003 FAB survey data.

Table 5
Net Worth Regression Results

Variable	Coeff.	Std. Err.	Pr > |t|
Expected-ideal gap	-0.130**	0.055	0.017
Ideal level	-0.067	0.044	0.127
Log 1999 income	0.135	0.170	0.425
Zero 1999 income	1.264*	0.728	0.082
Past income	0.509***	0.158	0.001
Zero past income	1.467**	0.695	0.035
Future income	-0.011	0.106	0.916
Zero future income	-0.037	0.454	0.934
Age	0.212***	0.045	0.000
Age^2	-0.001***	0.000	0.003
Empl. status			
Working		*Omitted*	
Partially retired	0.026	0.220	0.906
Retired	0.292	0.258	0.258
Occupation			
Faculty		*Omitted*	
Mgmt./Sen. Admin.	-0.156	0.153	0.308
Tech./Professional	0.022	0.146	0.881
Other	-0.099	0.170	0.559
Education			
College or below	-0.305**	0.143	0.033
M.A./Profesional		*Omitted*	
Ph.D.	-0.032	0.126	0.798
R. has DB plan	-0.203	0.126	0.108
S. has DB plan	-0.080	0.153	0.603
Marital status			
Curr. married		*Omitted*	
Prev. married	-0.600***	0.168	0.000
Never married	-0.344**	0.157	0.029
Male respondent	-0.047	0.111	0.673
Num. kids	0.028	0.062	0.650
Constant	-3.046***	1.117	0.006

Source: Authors' tabulation of 2003 survey data.
Notes: The dependent variable is log of net worth. We used a censored regression (Tobit) technique to include people with net worth of zero or less, as well as a multiple imputation process for censored values of the expected-ideal gap; see text. There were 374 observations used in this regression.

Table 6
Regressions for Wealth Categories

	Non-Retirement Assets			Retirement Assets		
Variable	Coeff.	S.E.	Pr $> \|t\|$	Coeff.	S.E.	Pr $> \|t\|$
Expected-ideal gap	-0.208**	0.087	0.017	-0.079	0.062	0.197
Ideal level	-0.075	0.076	0.329	-0.014	0.054	0.791
Log 1999 income	0.025	0.309	0.935	0.077	0.216	0.720
Zero 1999 income	1.069	1.623	0.510	1.387	1.134	0.221
Past income	0.867***	0.303	0.004	0.546***	0.212	0.010
Zero past income	3.495**	1.766	0.048	1.340	1.232	0.277
Future income	-0.024	0.184	0.894	-0.049	0.128	0.704
Zero future income	0.373	0.807	0.644	-0.106	0.563	0.851
Age	-0.106	0.101	0.294	0.283***	0.071	0.000
Age^2	0.001	0.001	0.199	-0.002***	0.001	0.004
Empl. status						
Working		*Omitted*			*Omitted*	
Partially retired	-0.268	0.387	0.489	0.424	0.270	0.116
Retired	-0.334	0.515	0.516	-0.044	0.360	0.903
Occupation						
Faculty		*Omitted*			*Omitted*	
Mgmt./Sen. Admin.	0.134	0.262	0.610	-0.087	0.183	0.634
Tech./Professional	0.010	0.254	0.968	0.068	0.178	0.701
Other	-0.003	0.304	0.992	-0.306	0.212	0.149
Education						
College or below	-0.500**	0.249	0.044	-0.368**	0.173	0.034
M.A./Profesional		*Omitted*			*Omitted*	
Ph.D.	0.356	0.222	0.108	-0.095	0.155	0.537
R. has DB plan	0.002	0.224	0.994	-0.260*	0.157	0.098
S. has DB plan	0.129	0.272	0.635	-0.025	0.190	0.894
Marital status						
Curr. married		*Omitted*			*Omitted*	
Prev. married	-0.204	0.294	0.487	-0.541***	0.205	0.008
Never married	-0.482*	0.278	0.083	-0.337*	0.194	0.083
Male respondent	-0.126	0.192	0.510	0.208	0.134	0.121
Num. kids	-0.065	0.107	0.543	0.004	0.074	0.962
Constant	2.208	2.278	0.332	-5.381***	1.593	0.001
N	362			362		

Source: Authors' calculations based on 2000, 2001, and & 2003 survey data.
Note: Dependent variables are natural logarithms of the quantities listed at head of each set of columns. Asterisks indicate the level of statistical confidence for rejection of the hypothesis that the relevant coefficient is (independently) equal to zero: "***" indicates rejection at better than a 1% level of confidence, "**" indicates rejection at better than a 5% level, and "*" indicates rejection at better than a 10% level.

Table7
The Revealed Preference Gap

RP Gap	(n)	(%)
10	1	0.1
5	16	1.2
4	5	0.4
3	14	1.1
2	17	1.3
1	44	3.3
0	1,173	88.3
-1	32	2.4
-2	19	1.4
-3	3	0.2
-4	3	0.2
-5	1	0.1
Total	1,328	100.0

Source: Authors' tabulations of 2003 FAB survey data.

Table 8
Probit Regression for Non-Zero EI Gap

Variable	Coeff.	Std. err.	Pr> $\|t\|$
Age	-0.016***	0.003	0.000
Male	-0.166**	0.764	0.030
Not Dependable	0.075***	0.035	0.032
Not Organized	0.154***	0.035	0.000
Ideal	-0.310***	0.32	0.000
Constant	1.657***	0.262	0.000

Source: Authors' tabulations of 2003 Survey Data