

Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving

Richard H. Thaler

Shlomo Benartzi

August 2001

Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving

Abstract

As firms switch from defined benefit plans to defined contribution plans, employees bear more responsibility for making decisions about how much to save. The employees who fail to join the plan, or who participate at a very low level, appear to be saving at less than the predicted life-cycle savings rates. Behavioral explanations for this behavior stress bounded rationality and self-control and suggest that at least some of the low-saving households are making a mistake, and would welcome an aid to help their saving decision making. In this paper we propose such a prescriptive savings plan, called Save More Tomorrow (hereafter, the SMT plan). The essence of the plan is straightforward: people commit in advance to allocate a portion of their future salary increases toward retirement savings. We also report evidence on the first implementation of the SMT plan. Our key findings are the following: (1) Most people (78 percent) who were offered the SMT plan elected to use it; (2) virtually everyone (98 percent) who joined the plan remained in it through two pay raises, and the vast majority (80 percent) remained in it through the third pay raise; and (3) The average saving rates for SMT plan participants increased from 3.5 percent to 11.6 percent over the course of 28 months. The results suggest that behavioral economics can be used to design effective prescriptive programs for important economic decisions.

Richard H. Thaler*
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637
TEL: (773) 702-5208
E-mail thaler@gsb.uchicago.edu

Shlomo Benartzi*
The Anderson School at UCLA
110 Westwood Plaza
Los Angeles, CA 90095-1481
TEL: (310) 206-9939; FAX: (310) 825-3165
E-mail: shlomo.benartzi@anderson.ucla.edu

* We are grateful to Brian Tarbox for implementing the Save More Tomorrow plan and for sharing with us the data. We are also grateful to Hewitt Associates and John Hancock for data and Financial Engines for access to their software. Thanks also to Ted O'Donoghue, David Laibson, Casey Mulligan, and Cass Sunstein for helpful comments.

1. Introduction

Economic theory generally assumes that people solve important problems as economists would. The life-cycle theory of saving is a good example. Households are assumed to want to smooth consumption over the life-cycle and are expected to solve the relevant optimization problem in each period before deciding how much to consume and how much to save. Actual household behavior might differ from this optimal plan for either of two reasons. First, the problem is a hard one, even for an economist, so households might fail to compute the correct saving rate. Second, even if the correct savings rate were known, households might lack the self-control to delay current consumption in favor of future consumption (Thaler and Shefrin, 1981).

One fact that underscores the important role of self-control is that the typical American household accumulates retirement wealth primarily in three forms: social security, pensions, and home equity. Neither social security nor defined-benefit pension plans require willpower on the part of participants, and once a home is purchased, the monthly mortgage bill provides a useful discipline in building up equity.

Those Americans who have access to and make use of all three low-willpower savings techniques appear to be doing a decent job of saving for retirement. Gustman and Steinmeir (1998), using the 1992 Health and Retirement Survey of households with heads born between 1931 and 1941, find that households with pensions have what appear to be adequate income replacement rates. A majority of the pensions in their sample are of the defined-benefit (DB) variety, however, in which self-control plays no role. Over the past decade, there has been a rapid change toward defined-contribution (DC) plans that require employees to actively join and select their own savings rate. For those workers who are only eligible for a DC plan and elect not to join or to contribute a token amount, savings adequacy may be much lower. One hint at this comes from Gustman and Steinmeir's analysis of workers who do not have pensions. Their wealth and savings adequacy levels are substantially lower than those with pensions. Indeed, those workers with pensions are wealthier by approximately the value of their pension.¹

For whatever reason, some employees at firms that only offer defined-contribution plans contribute little or nothing to the plan. In this paper we take seriously the possibility that some of these low saving workers are making a mistake. By calling their low-saving behavior a mistake, we mean that they might characterize the action the same way, just as someone who is 100 pounds overweight might agree that he or she weighs too much. We then use principles from psychology and behavioral economics to devise a program to help people save more. The program is called Save More Tomorrow (SMT), and the basic idea is to give workers the option of committing themselves now to increase their savings rate later, each time they get a raise. We report data from one firm that has implemented the program.

¹ It is sometimes argued that this fact can be explained by selection effects (those workers with a "taste for savings" go to work for companies with more attractive pension benefits), but it is important not to push this argument too far. It is implausible that pension benefits are so salient that workers sort themselves to firms primarily on this feature. Many other features of a job determine its attractiveness, and potential employees must make tradeoffs. (To give one example, one of the authors of this paper is much more interested in collegiate athletics than the other, but he teaches at the University of Chicago, not UCLA!) Therefore, we should not expect underlying preferences and employment characteristics to be perfectly correlated on any single dimension.

We note that the null hypothesis predicted by the standard economic approach is that workers will have no interest in joining the SMT plan. If households are already choosing their optimal life-cycle savings rate, then they will not join a program that will commit them to periodic changes. In contrast, the behavioral economics prediction is that workers will find this program quite attractive and that it will significantly increase the savings rates of those who join the plan.

2. A Prescriptive Approach to Increasing Savings Rates

Howard Raiffa (1982) suggested that economists and other social scientists could benefit from distinguishing three different kinds of analyses: normative, descriptive, and prescriptive. Normative theories characterize rational choice and are often derived by solving some kind of optimization problem. The life-cycle hypothesis is an example of a normative theory of saving since it is based on the solution to a lifetime consumption-smoothing problem. Descriptive theories simply model how people actually choose, often by stressing systematic departures from the normative theory. In the realm of savings behavior, Shefrin and Thaler (1988) offer the behavioral life-cycle hypothesis as a descriptive model of household savings in which self-control and mental accounting play key roles. Finally, prescriptive theories are attempts to offer advice on how people can improve their decision making and get closer to the normative ideal. Prescriptions often have a second-best quality. For a golfer who hits a slice (in which the ball tails off to the right) when he would prefer to hit the ball straight, simple prescriptive advice might be to aim to the left. Better prescriptive advice would help the golfer hit the ball straight. This paper is an attempt at good prescriptive savings advice.

Before writing a prescription, one must know the symptoms of the disease being treated. Households may save less than the life-cycle rate for various reasons. First, determining the appropriate savings rate is difficult, even for someone with economics training. Since the switch from DB to DC savings plans is recent, there are not yet satisfactory heuristics that approximate a good solution to the problem.² One obvious solution to this problem is financial education (Bernheim, Garrett, and Maki, 1997). Second, saving for retirement requires self-control. When surveyed about their low savings rates, many households report that they would like to save more but lack the willpower. For example, Choi, Laibson, Madrian and Metrick (2001) report that two-thirds of their sample of 401(k) participants think their savings rate is “too low”.³ A third problem, closely related to self-control, is procrastination, the familiar tendency to postpone unpleasant tasks. In Choi et al’s group of self-reported undersavers, 35 percent express an intention to increase their savings rate in the next few months, but 86 percent of these well-intended savers have made no changes to their plan four months later.

Self-control and procrastination used to be strange concepts to economists, but are now topics of growing interest by behavioral economics theorists (e.g., Laibson, 1997; O’Donoghue

2 The most common heuristics in place appear to be to save the maximum allowed by law or to save the minimum necessary to receive the full “match” offered by the employer. Neither of these amounts were computed to be solutions to the life-cycle savings problem. In any case, the group our program is aimed at are not using either of these heuristics.

3 Similarly, a 1997 survey by Public Agenda finds that 76 percent of respondents think they should be saving more for retirement. See Farkus and Johnson (1997) for details.

and Rabin, 1999). Modern models of these problems use the concept of hyperbolic discounting (see Ainslie, 1975). Since Strotz's (1955) early paper, economists have known that intertemporal choices are only time consistent if agents discount exponentially using a discount rate that is constant over time. But there is considerable evidence that people display time-inconsistent behavior, specifically, weighing current and near-term consumption especially heavily. Consider a choice between two rewards, a small one at time t (S_t) and a big one at time $t+1$ (B_{t+1}). When t is far off, agents prefer B_{t+1} , as the difference in the value of the prizes exceeds the perceived costs of waiting. But as t approaches 0, the ratio of discounted values increases, causing people to switch their preferences.⁴ Such present-biased preferences can be captured with models that employ hyperbolic discounting. These models come in two varieties: sophisticated and naïve. Sophisticated agents (modeled by Laibson) realize they have hyperbolic preferences and take steps to deal with the problem, whereas naïve agents fail to appreciate at least the extent of their problem (see O'Donoghue and Rabin, 1999, 2001). Actual behavior is likely best described by something between naiveté and sophistication.

Hyperbolic agents procrastinate because they (wrongly) think that whatever they will be doing later will not be as important as what they are doing now. The more naïve agents are, the more pronounced is the tendency to procrastinate. Procrastination, in turn, produces a strong tendency toward inertia or what Samuelson and Zeckhauser (1988) have dubbed status quo bias. Status quo bias is prevalent in the retirement saving domain. For example, Samuelson and Zeckhauser report on the behavior of the 1987 participants of TIAA-CREF, the large retirement plan that then catered to university employees. Their analysis reveals that the *median* number of changes in the asset allocation over the *lifetime* was zero! In other words, more than half the participants in TIAA-CREF reached retirement with the same asset allocation as the day they became eligible for the plan. Note that zero changes means that participants were electing a constant *flow* into the two funds then offered, TIAA, a bond fund, and CREF, a stock fund, and engaged in no rebalancing. Since stocks appreciated much more than bonds over this period, participants with a constant flow (such as 50-50, the most common allocation) ended up with a much larger share in stocks over time. A recent study by Ameriks and Zeldes (2000), using a ten-year panel of TIAA-CREF participants, finds a similar result. Nearly half of the participants made no changes to their plan over the ten-year period.⁵

The importance of procrastination and status quo bias in the design of prescriptive savings plans is illustrated by the experience some firms have had with so-called automatic enrollment plans. In such plans, when employees first become eligible for the saving plan they are automatically enrolled unless they explicitly opt out. So, unlike the typical plan, in which the default is not to join, here the default is to join. Employees who take no action are typically enrolled at a modest saving rate (such as 3 percent) and a conservative investment strategy. Standard economic theory would predict that this change would have virtually no effect on saving behavior. The reduction in the costs of joining the plan (typically filling out a short form) are trivial compared with the potential benefits of the tax-free accumulation of wealth, and in some cases a "match" is provided by the employer, in which the employer typically contributes

4 For evidence on hyperbolic discounting see Thaler (1981) and the papers in Loewenstein and Elster (1992).

5 Choi, Laibson and Metrick (2000) find somewhat more frequent trading in a sample of workers at two firms in 1998 and 1999, partly due to the ease of trading via the internet which was possible at both firms. However, this increase in trading may also be attributable to rapidly rising stock prices during this period, and the resulting excitement among individual investors.

50 cents to the plan for every dollar the employee contributes, up to some maximum. In contrast, if agents display procrastination and status quo bias, then automatic enrollment could be useful in increasing participation rates.

Consistent with the behavioral predictions, automatic enrollment plans have proven to be remarkably successful in increasing enrollments. In one plan studied by Madrian and Shea (1999), participation rates for newly eligible workers increased from 49 percent to 86 percent. Other plans have obtained participation rates over 90 percent. (Choi, Laibson, Madrian and Metrick, 2001b). But there is a down side to automatic enrollment. The very inertia that explains why automatic enrollment increases participation rates can also lower the saving rates of those who do participate. In the firm Madrian and Shea studied, the vast majority of new enrollees elected the default saving rate (3 percent), and Madrian and Shea's analysis shows that many of these employees would have elected a higher saving rate if left to their own devices. (Choi et al 2001b explore these issues in depth.) A goal of the SMT plan is to obtain some of the advantages of automatic enrollment while avoiding some of the disadvantages.

Based on our analysis of undersaving households in the previous section, some elements of a proposed solution are fairly obvious. The presence of bounded rationality suggests that the program should be simple, and should help people approximate the life-cycle saving rate if they are unable to do so themselves. Hyperbolic discounting implies that opportunities to save more in the future will be considered more attractive than those in the present. Procrastination and inertia suggest that once employees are enrolled in the program, they should remain in until they opt out.

The final behavioral factor that should be considered in designing a prescriptive savings plan is loss aversion, the empirically demonstrated tendency for people to weigh losses significantly more heavily than gains. Estimates of loss aversion typically hover around 2.0: losses hurt roughly twice as much as gains yield pleasure. These estimates come both from risky choice (Tversky and Kahneman, 1992) and from riskless choice (Kahneman, Knetsch, and Thaler, 1990).

Loss aversion affects savings because once households get used to a particular level of disposable income, they tend to view reductions in that level as a loss. Thus, households may be reluctant to increase their contributions to the savings plan because they do not want to experience this cut in take-home pay. Significantly, gains and losses appear to be experienced in nominal dollars. For example, in a study of perceptions of fairness (Kahneman, Knetch and Thaler, 1986), subjects were asked to judge the fairness of pay cuts and pay increases in a company located in a community with substantial unemployment. One group of subjects was told that there was no inflation in the community and was asked whether a 7 percent wage cut was "fair." A majority, 62 percent, judged the action to be unfair. Another group was told that there was 12 percent inflation and was asked to judge the perceived fairness of a 5 percent raise. Here, only 22 percent thought the action was unfair. Similar results suggesting this money illusion are reported by Shafir, Diamond and Tversky (1997). The combination of loss aversion and money illusion suggests that the time of pay increases may be a propitious time to try to get people to save more, since they are less likely to consider an increased contribution to the plan as a loss than at other times of year.

To summarize, for households that appear to be saving too little, the behavioral analysis stresses four factors that are important explanatory factors: bounded rationality, self-control, procrastination (which produces inertia), and nominal loss aversion. These households are not sure how much they should be saving, though they realize that it is probably more than they are

doing now, but they procrastinate about saving more now, thinking they will get to it later. Our program to increase saving is aimed at this group.

3. The Save More Tomorrow Program

Our goal was to design a program to help those employees who would like to save more but lack the willpower to act on this desire. Based on the principles discussed so far, we have proposed a program we call Save More Tomorrow (SMT). The plan has the following ingredients: First, employees are approached about increasing their contribution rates a considerable time before their scheduled pay increase. Because of hyperbolic discounting, the lag between the sign-up and the start-up date should be as long as feasible.⁶ Second, if employees join, their contribution to the plan is increased beginning with the first paycheck after a raise. This feature mitigates the perceived loss aversion of a cut in take-home pay. Third, the contribution rate continues to increase on each scheduled raise until the contribution rate reaches a preset maximum. In this way, inertia and status quo bias work toward keeping people in the plan. Fourth, the employee can opt out of the plan at any time. Although we expect few employees to be unhappy with the plan, it is important that they can always opt out. Knowledge of this feature will also make employees more comfortable about joining.

The first implementation of the SMT plan took place in 1998 at a midsize manufacturing company.⁷ Prior to the adoption of the SMT plan, the company suffered from low participation rates as well as low saving rates. This was a concern for two reasons. First, since the company did not have a defined-benefit plan, management was concerned that some of the workers might not be saving enough to support themselves when they retired. Second, the company was being constrained by U.S. Department of Labor anti-discrimination rules that restrict the proportion of benefits that can be paid to the higher-paid employees in the firm. Since the lower paid workers were the ones who were typically saving little or nothing, the executives were not able to contribute the maximum normally allowed to their own plan.

We did not have control over the exact implementation of the SMT plan in this company. The company, with the help of an investment consultant, decided the specific details of how the plan would work, and their choices are somewhat idiosyncratic. Nevertheless, we think much can be learned from this company's experience. As other companies adopt the SMT plan, we will gain information on the relative merits of specific implementation decisions.

In an effort to increase the savings rates of the employees, the company hired an investment consultant and offered his services to every employee eligible for its retirement savings plan. Of the 315 eligible participants, all but 29 agreed to meet with the consultant and get his advice. Based on information that the employee provided, the consultant used commercial software to compute a desired saving rate, which can be thought of as an estimate of the appropriate life-cycle savings rate. The consultant also discussed with each employee how much of an increase in savings would be considered economically feasible. If the employee seemed very reluctant to increase his or her saving rate substantially, the consultant would

⁶ The intuition here is the same as why requests to give a talk or write a chapter meet with more success when they are received many months ahead of time.

⁷ The company prefers to remain anonymous.

constrain the program to increase the saving contribution by no more than 5 percent.⁸ The consultant justified his decision not to go with the advice from the program mechanically as follows: “In most cases with rank and file workers, the computer program calculates that workers contribute the maximum [allowed by the IRS and the plan rules] and makes that recommendation. As a practical matter, when the average worker receives this recommendation from the computer program or the "financial planner," s/he shuts down and does nothing. So in all cases, after we reviewed their current plan but before I hit the 'Get Advice' button, I would discuss willingness to save with each participant. As you can imagine, the majority of workers live paycheck to paycheck and can barely make ends meet, and they tell you that immediately. ... If a participant indicated a willingness to immediately increase their deferral level by more than 5 percent, I hit the 'Get Advice' button. Otherwise, I would constrain the advice proposed to an increase of no more than 5 percent.”

Of the 286 employees who talked to the investment consultant, only 79 (28 percent) were willing to accept the consultant’s advice, even with this adjustment to constrain some of the increases to 5 percent increase in their saving rates. For the rest of the participants, the planner offered a version of the SMT plan, proposing that they increase their saving rates by 3 percentage points a year starting with the next pay increase. This was quite aggressive advice, since pay increases were barely more than this amount (averaging 3.25 percent for hourly employees and 3.47 percent for salaried employees). The pay increases were scheduled to occur roughly three months from the time the advice was being given. With the 3 percent a year increases, employees would typically reach the maximum tax deferred contribution within four years.

Even with this aggressive strategy of increasing saving rates, the SMT plan proved to be extremely popular with the participants. Of the 207 participants who were unwilling to accept the saving rate proposed by the investment consultant, 162 (78 percent) agreed to join the SMT plan. More importantly, the majority of these participants did not change their mind once the savings increases took place. Only 4 participants (2 percent) dropped out of the plan prior to the second pay raise, with 29 more (18 percent) dropping out between the second and third pay raises.⁹ Hence, the vast majority of the participants (80 percent) have remained in the plan through three pay raises. Furthermore, even those who withdrew from the plan did not reduce their contribution rates to the original levels; they merely stopped the future increases from taking place. So, even these workers are saving significantly more than they were before joining the plan. Table 1 provides data on participation in the plan.

[Insert Table 1 About Here]

The impact of the SMT plan on saving is shown in Table 2. All data refer only to those employees who have remained with the company for the entire period of study. When the investment consultant was introduced, overall saving in the plan was 4.4 percent. The employees who did not want to talk to the investment consultant were saving more than the average, 6.6 percent. The group that accepted the advice of the consultant started at exactly the

⁸ Here and elsewhere, when we mention an increase in the saving rate by some percentage amount we are referring to an increase of percentage points, e.g., from a 2% saving rate to a 7% saving rate.

⁹ Interestingly, most of the employees who dropped out between the second and third pay increases worked for a single supervisor who apparently disapproved of the SMT plan.

group average, 4.4 percent, which was raised to 9.1 percent. At the end of our data collection period, that had slipped slightly to 8.7 percent. Those who were unwilling to accept the advice were, not surprisingly, starting from a lower base of 3.5 percent and, so, would find the advice harder to adopt. Once they got their first pay raise, however, their saving rate jumped to 6.5 percent, and after their second and third pay raises it was up to 9.4 percent and 11.6 percent, respectively. In fact, those participating in the SMT plan ended up with a higher saving rate than those who accepted the consultant's recommendation.

[Insert Table 2 About Here]

Of course, the implementation of the SMT plan was not conducted as an experiment with random assignment to conditions. Participants selected themselves into the SMT plan. In other circumstances, one might worry that the observed increase in savings rates might be attributable to some unmeasured “taste for savings” in the households that join the SMT plan, however, this worry seems inappropriate here on two counts. First, the SMT participants had been saving very little before joining the plan, so one would have to believe that their taste for saving was newly acquired. Second, recall that the SMT plan was only offered to those employees who were unwilling to increase their savings rate immediately by 5 percent. So, if anything, the group that accepted the consultant’s advice would appear to have a greater taste for saving than those in the SMT plan.

The design of our study also rules out another potential explanation for our results based on information. Since the employees met with the investment consultant, it is true that they received useful information about proper savings rates and this information could sensibly affect their savings rates. However, all the employees that agreed to meet with the consultant received this information, including those who accepted the consultant’s advice to increase their savings rate immediately. We find it difficult to construct an information-based explanation for the subsequent increases in savings rates for those enrolled in the SMT plan.

4. How Successful Is SMT as a Prescriptive Device?

Obviously, the SMT plan has a dramatic effect on saving rates. Those who joined the plan more than tripled their saving rates in 28 months. This raises the question of what effect SMT has on savings adequacy. Is this increase enough to make a substantial difference in the standard of living the participants will have in retirement? If so, is it possible that we have been “too successful” and have somehow duped the participants into saving too much. This section offers some information on these important questions.

We do not have demographic or financial information about the employees in our study, so we need to make some assumptions about their household financial situation in order to calculate the likely effects of joining the SMT plan. We make calculations for workers who join the plan at age 25, 35, 45, or 55, for three different annual incomes: \$25,000, \$50,000 and \$75,000. We estimate beginning 401(k) account balances, using data from Hewitt Associates, of some of the larger 401(k) plans they administer. In particular, we calculate the account balances of people of a similar age, income, and saving rates. To avoid the issue of multiple 401(k) accounts per individual, we select only those who remained with the same employer through their career. As to savings and investment choices, we assume employees are saving 4 percent in the 401(k) plan when they join the SMT plan, which increases to 12 percent. We also assume that the employer matches employee contributions at a 50 percent rate on the first 6 percent of

employee contributions (as was true in the firm we studied). For other financial assets we assume that non-401(k) employee savings are half the existing balance in the 401(k) account, based on data from John Hancock Financial Services (1999). Finally, we assume that employees choose a portfolio mix of 60 percent stocks and 40 percent bonds. The particular company in our study does not sponsor a defined-benefit pension, so we assumed no pension benefits. Last, we assumed the statutory benefits from Social Security. When these data are entered into Financial Engines, the software provides several points on a probability distribution of retirement income, as shown in Table 3. We report the 50th percentile of this distribution.

[Insert Table 3 About Here]

Table 3 reports retirement income replacement rates, that is, the ratio of retirement income to pre-retirement income. There is not a consensus among economists on the appropriate replacement rate (see Bernheim (1993), Boskin and Shoven (1987), and Gustman and Steinmeier (1998) for discussions of this issue), but roughly speaking replacement rates near 100 percent seem adequate while numbers below 70 percent many consider to be too low.¹⁰ Panel A shows the expected income replacement rates for our employees before they join the SMT plan, all of which are 60 percent or less. Replacement rates are highest for the \$25,000 income category because social security offers substantial replacement at that level. Panel B shows that replacement income rates are considerably higher with the SMT plan, especially for those joining the plan when young. Obviously, increasing the savings rate is less effective when starting at 55 than at 25. Still, expected replacement rates are never above 100%, so there does not appear to be a problem that we have induced people to save too much. Furthermore, if the stock market returns are exceptionally high, workers can always reduce savings rates as they approach retirement and have a better idea of what their retirement income will be.

5. The Potential Effect of SMT on the US Personal Savings Rate

The first implementation of the SMT program was extremely successful, and we have shown that the SMT program can be an effective way of increasing the savings rates of some employees. It is natural to ask what effect the program might have on personal savings in the U.S. if the program were adopted widely. In this section, we explore the potential effect of the SMT program on other retirement saving plans by using data from Hewitt Associates. The data set includes demographic and account balance information on the participants in 15 large companies, covering a total of 539,516 employees. Based on comparisons with data from Fidelity (2001) and John Hancock Financial Services (1999), two other large 401(k) service providers, we believe our sample is representative at least of employees of large companies. Consequently, we think it can serve as a basis for some rough estimates on the potential contribution SMT can make in increasing employee savings rates.

¹⁰ One might think that a 100 percent replacement rate would be too high, suggesting very that agents are very patient. However, survey evidence suggests that households desire an increasing consumption profile, even if interest rates are very low. Laibson (1999) offers a cogent discussion of this issue, and also reports that economists also prefer rising profiles for themselves. If agents want a rising profile then even a 100 percent replacement rate may be too low.

We have made projections on three different hypothetical implementation strategies. For all three plans we assume that the saving rate is set to increase by either 1, 2, or 3 percent per year, and we assume that 5 percent of people drop out of the program each year, leaving their savings rate at the level they had obtained up to that point. The differences in the assumptions are based on the assumed success rates of getting employees to enroll. Plan A is based on the actual plan discussed above, which used one-on-one interactions with a financial consultant. Based on our experiences, we assume that 80 percent of those who are currently participating in the savings program will join the SMT plan, and half of those who are not enrolled will join.

Plan B is based on the assumption that the SMT plan is only marketed to employees with a direct mail campaign rather than personal contact. This approach is much less costly but is also less effective in reaching potential enrollees. Based on a recent experience we have had using this method in another company, we project 20 percent enrollment rates for those currently in the savings plan and 10 percent for those who are not currently saving anything.

Plan C is to combine the SMT program with automatic enrollment. Employees would be enrolled in the SMT plan unless they opted out. Based on our experience and those of Madrian and Shea (1999) and Choi et al (2001b), for Plan C we estimate that 90 percent of the employees would join the program. The saving rates we report are weighted by income, and they are averaged across all employees (whether or not they are saving). Hence, the reported rates represent the average savings per dollar of income.¹¹

[Insert Table 4 About Here]

The results of our projections are displayed in Table 4. As of year-end 2000, the saving rate in the Hewitt sample averaged 5.0 percent, less than a third of the allowable IRS deferrals, which averaged 17.7 percent for our sample.¹² This means that there is considerable opportunity for the SMT program to increase the saving rate. With Plan A, using one-on-one interaction with a financial consultant, the SMT program could boost the saving rate from 5.0 percent to 9.7 percent within five years (see Panel A). Switching to the cheaper method of one-shot mailing, the effects are much smaller (see Panel B). For instance, over the course of five years, the saving rate would increase from 5.0 percent to 6.1 percent. But if employees were automatically enrolled in the program, as in Plan C, the average saving rate is projected to increase from 5.0 percent to 10.9 percent within five years (see Panel C).

How large is the potential increase in saving rates? In terms of dollars, we calculate that each one-percentage-point increase in the employee saving rate would translate into \$250 million of additional annual contributions for the Hewitt sample. Extrapolating from our sample of half a million individuals to the universe of 50-60 million individuals with access to 401(k) plans, we estimate roughly \$25 billion of additional annual contributions for each 1 percent increase, so if a 5 percent increase were obtained, this would increase personal saving by \$125 billion per year. Percentage-wise, this would amount to 1.5 percent of disposable personal income (Bureau of Economic Analysis, 2000). Since the current personal savings rate is hovering near zero, this is

¹¹ For simplicity, our calculations exclude the effects of employer contributions and employee turnover. On one hand, employer contributions will probably increase the effect of the SMT program, because increased employee contributions will often trigger higher employer contributions. On the other hand, employee turnover is likely to decrease the effect of the SMT program.

¹² The IRS limit for the year 2000 was the lower of \$10,500 and 25 percent of income.

a substantial increase. Furthermore, unlike other approaches to increasing the employee saving rate, such as increasing the maximum allowable contribution, much of the gains from the SMT program come from those who are saving little or nothing now. This means that the increase can be presumed to be virtually all “new” savings as opposed to substitution from other (possibly taxable) forms.

6. Conclusions

The initial experience with the SMT plan has been remarkably successful. Most of the people who were offered the plan elected to use it. Furthermore, the majority of the people who joined the SMT plan stick with it. Consequently, SMT participants more than tripled their saving rates. Of course, the reason why the SMT plan works so well is that inertia is so powerful. Once people enroll in the plan, few will ever get around to opting out. The SMT plan takes precisely the same behavioral tendency that induces people to postpone saving indefinitely (i.e., procrastination and inertia) and puts it to use. As the consultant has pointed out to us, there is an argument for offering the SMT plan to all participants, even those who would have been willing to make their initial savings increase more than the first step of the SMT plan (here 3 percent). Since even these eager savers never got around to changing their savings allocations again the following year, the SMT plan participants were already saving more after just 16 months (see Table 2).

One objection we sometimes hear about both the SMT plan and automatic enrollment is that they amount to some kind of sneaky paternalism. We think such criticism is misplaced. Notice that no one is being forced to do anything here. The whole point of choosing a default rule rather than a mandatory rule is to give people choices. As to the choice of the specific default rule, there is no avoiding having some default rule. Policy makers always have to choose what happens to people who fail to make a choice. Making the default the one that the policy makers believe is the one most people would choose upon reflection is intended to be helpful. For example, such a choice about the default option is involved when an employee is allowed by law to elect to make some payroll deduction (e.g., parking or health insurance) on a pretax basis. If most employees choose to do this, and it is an option for everyone, then courteous employers make this the default, not to be intrusive but to help absent-minded employees, who may forget to return the correct form and therefore lose the benefit. Similarly, we do not think that automatic enrollment is intrusive, and the SMT plan as implemented in this case study is even less intrusive since it required the worker to actively join the plan. Of course, the SMT plan works precisely because the default rule is to stay in the plan once having joined. And although one might argue that workers enrolled in the SMT plan will end up saving more than they will need, our calculations indicate that this is probably not the case.

Finally, we hope this study serves as a valid reply to two frequent critiques of behavioral economics: the reliance on laboratory studies using modest stakes; and the ex-post explanation of anomalous facts drawing on what is alleged to be an limitless store of potential behavioral explanations. Here, we have used behavioral principles to design a plan to increase savings rates and tested the idea in the real world.

References

- Ainslie, G., 1975, Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control,' *Psychological Bulletin* 82.4, 463-496.
- Ameriks, J., and S. Zeldes, 2000, How do Household Portfolio Shares Vary with Age? working paper, Columbia University.
- Bernheim, B. D., 1993, *Is the Baby Boom Generation Preparing Adequately for Retirement?* Merrill Lynch & Co.
- Bernheim B. D., D. M. Garrett, and D. M. Maki, 1997, Education and Saving: The Long-Term Effects of High School Financial Curriculum Mandates, NBER Working Paper No. W6085.
- Boskin, M. J., and J. Shoven, 1987, Concepts and Measures of Earnings Replacement Rates During Retirement. In Bodie, Z., J. Shoven, and D. Wise, editors, *Pensions and Retirement in the United States*. Chicago: University of Chicago Press for NBER. 113-141.
- Bureau of Economic Analysis, 2000, Data from the National Income and Product Accounts: Personal Income and its Disposition, <http://www.bea.doc.gov/briefrm/tables/abr6.htm>.
- Choi, D, D. Laibson, B. Madrian, and A. Metrick, 2001, For Better or For Worse: Default Effects and 401(k) Savings Behavior. Forthcoming in D. Wise ed. *Frontiers in the Economics of Aging*, NBER, 2001.
- Choi, D, D. Laibson, and A. Metrick, 2001, Does the Internet Increase Trading? Evidence from Investor Behavior in 401(k) Plans, NBER Working Paper #7878.
- Farkus, S. and J. Johnson, 1997, *Miles to Go: A Status Report on Americans Plans for Retirement*. New York: Public Agenda
- Fidelity Investments, 2001, A Report on Corporate Defined Contribution Plans.
- Gustman, A. L., and T. L. Steinmeir, 1998, Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study, NBER Working Paper 6681.
- John Hancock Financial Services, 1999, The sixth defined contribution plan survey.
- Kahneman, D., J. L. Knetsch, and R. H. Thaler, 1990, Experimental Tests of the Endowment Effect and the Coase Theorem, *Journal of Political Economy* 98.6, 1325-1348.
- _____, 1986, Fairness as a Constraint on Profit Seeking, *American Economics Review* 76, 728-741.
- Laibson, D., 1997, Golden Eggs and Hyperbolic Discounting, *Quarterly Journal of Economics* 112, 443-477.
- Laibson, D. 1999, Discussion of Engen, Gale and Uccello, *Brookings Papers on Economic Activity*. 174-177.
- Loewenstein, G. and J. Elster (eds), 1992, Choice Over Time, New York: Russell Sage.
- Madrian, B. C., and D. Shea, 1999, The Power of Suggestion: An Analysis of 401(k) Participation and Saving Behavior, Working Paper, University of Chicago.
- O'Donoghue, T., and M. Rabin, 1999, Doing It Now or Later, *American Economic Review* 89.1, 103-124.
- O'Donoghue, T., and M. Rabin, 2001, *Quarterly Journal of Economics*, 116.1, 121-160.
- Raiffa, H. (1982), *The Art and Science of Negotiation*. (Harvard University Press: Cambridge MA)

- Samuelson, W., and R. J. Zeckhauser, 1988, Status Quo Bias in Decision Making, *Journal of Risk and Uncertainty* 1, 7-59.
- Shafir, E, P. Diamond and A. Tversky, 1997, Money Illusion, *Quarterly Journal of Economics* CXII (May), 341-374.
- Shefrin, H. M., and R. H. Thaler, 1988, The Behavioral Life-Cycle Hypothesis, *Economic Inquiry*, 26.4, 609-643.
- Strotz, R. H., 1955, Myopia and Inconsistency in Dynamic Utility Maximization, *Review of Economic Studies* 23.3, 165-180.
- Thaler, R. H., 1981, Some Empirical Evidence on Dynamic Inconsistency," *Economic Letters* 8, 201-207.
- Thaler, R. H., and H. M. Shefrin, 1981, An Economic Theory of Self-Control, *Journal of Political Economy* 89, 392-410.
- Tversky, A., and D. Kahneman, 1992, Advances in Prospect Theory: Cumulative Representation of Uncertainty, *Journal of Risk and Uncertainty* 5, 297-323.

Table 1
Participation Data

The number of plan participants prior to the adoption of the SMT plan	315
The number of plan participants who elected to receive a recommendation from the consultant	286
The number of plan participants who implemented the consultant's recommended saving rate	79
The number of plan participants who were offered the SMT plan as an alternative	207
The number of plan participants who accepted the SMT plan	162
The number of plan participants who opted out of the SMT plan between 1 st and 2 nd pay raise	4
The number of plan participants who opted out of the SMT plan between 2 nd and 3 rd pay raise	29
Overall participation rate prior to the advice	64%
Overall participation rate shortly after the advice	81%

Table 2
Average Saving Rates

	Participants who did not contact the financial consultant	Participants who accepted the consultant's recommended saving rate	Participants who jointed the SMT plan	Participants who declined the SMT plan	All
Number of participants	29	79	162	45	315
Pre-advice	6.6%	4.4%	3.5%	6.1%	4.4%
1 st Pay Raise	6.5%	9.1%	6.5%	6.3%	7.1%
2 nd Pay Raise	6.8%	8.9%	9.4%	6.2%	8.6%
3 rd Pay Raise	6.6%	8.7%	11.6%	6.1%	9.8%

Participants were offered access to an investment consultant. Many of those contacting the consultant found the recommended saving rates too high, and they were offered the SMT plan as an alternative.

Table 3
Median Income Replacement Ratios

Income	Age			
	25	35	45	55
<i>Panel A: Pre-SMT</i>				
\$25,000	60%	58%	54%	50%
\$50,000	51	50	50	52
\$75,000	47	45	46	41
<i>Panel B: Post-SMT</i>				
\$25,000	93%	81%	68%	56%
\$50,000	84	72	63	58
\$75,000	80	68	60	47

The above table displays the median income replacement ratios for different age and income profiles, using investment advice software by Financial Engines. The projections are based on the following assumptions: no defined-benefit pension, statutory Social Security benefits, employee saving rate of four percent before SMT and 12 percent thereafter, employer match of fifty cents on the dollar up to six percent, portfolio mix of 60 percent stocks and 40 percent bonds, and retirement age of 65.

Table 4
Projected Saving Rates

Projected Saving Rates with SMT in Year:							
SMT Annual Increments	0	1	2	3	4	5	10
<i>Panel A: One-on-One Interaction with a Financial Consultant</i>							
1%	5.0%	5.6%	6.2%	6.7%	7.2%	7.6%	9.2%
2	5.0	6.2	7.3	8.2	9.0	9.7	11.9
3	5.0	6.8	8.3	9.5	10.6	11.4	12.9
<i>Panel B: One-Shot Mailing</i>							
1	5.0	5.2	5.3	5.4	5.5	5.6	6.0
2	5.0	5.3	5.6	5.8	6.0	6.1	6.7
3	5.0	5.4	5.8	6.1	6.3	6.5	6.9
<i>Panel C: Automatic Enrollment</i>							
1	5.0	5.8	6.4	7.1	7.7	8.2	10.2
2	5.0	6.5	7.8	8.9	10.0	10.9	13.7
3	5.0	7.2	9.0	10.6	11.9	13.0	15.0

This table displays projected saving rates with SMT. Pre-SMT saving rates are based on actual data from Hewitt Associates, covering 539,516 employees. The effect of the SMT program is projected from our preliminary experience with the program. In Panel A, we used our first experience with the program, in which employees had one-on-one interaction with a financial consultant. Specifically, we assumed that the likelihood of current savers to join the program is 0.8 and the likelihood of non-savers to join is 0.5. In Panel B, we utilized our second experience with the program, in which employees were invited to join the program by mail. Here, we assumed that the likelihood of current savers to join the program is 0.2 and the likelihood of non-savers to join is 0.1. In Panel C, we assumed that employees are automatically enrolled in the program, unless they opt out. Here, the likelihood of joining the program is set to 0.9. The likelihood of subsequent opting out is 0.05 per year in all panels. The reported saving rates are weighted by income, and they are averaged across all employees (whether or not they are saving).