

Save More Tomorrow[™]: Using Behavioral Economics to Increase Employee Saving

Richard H. Thaler

University of Chicago

Shlomo Benartzi

University of California, Los Angeles

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 aid in making decisions about their saving. In this paper, we propose such a prescriptive savings program, called Save More Tomorrow[™] (hereafter, the SMarT program). The essence of the program is straightforward: people commit in advance to allocating a portion of their future salary increases toward retirement savings. We report evidence on the first three implementations of the SMarT program. Our key findings, from the first implementation, which has

We are grateful to Brian Tarbox for implementing the Save More Tomorrow[™] plan and for sharing the data with us. We would also like to thank many people at the following companies for their help: Financial Engines, Hewitt Associates, Ispat Inland, John Hancock, Philips Electronics, and the Vanguard Group. Jodi Dicenzo, Bill Sharpe, and Steve Utkus deserve special thanks. We are also grateful for comments from David Laibson, Brigitte Madrian, Casey Mulligan, Ted O'Donoghue, and Cass Sunstein. Benartzi would like to thank Reish Luftman McDaniel & Reicher for financial support. Save More Tomorrow is a registered trademark of Benartzi and Thaler, but the plan is available at no charge to any company that is willing to share data on the outcomes. This paper is dedicated to Sherwin Rosen, Thaler's thesis advisor. Thaler would not be an economist today if not for Rosen's help. The usual disclaimer, assigning none of the blame for errors to those thanked above, applies in spades to Sherwin. He would not have liked this paper much, but we sure would have enjoyed hearing him complain about it!

[*Journal of Political Economy*, 2004, vol. 112, no. 1, pt. 2]
© 2004 by The University of Chicago. All rights reserved. 0022-3808/2004/11201S1-0018\$10.00

been in place for four annual raises, are as follows: (1) a high proportion (78 percent) of those offered the plan joined, (2) the vast majority of those enrolled in the SMarT plan (80 percent) remained in it through the fourth pay raise, and (3) the average saving rates for SMarT program participants increased from 3.5 percent to 13.6 percent over the course of 40 months. The results suggest that behavioral economics can be used to design effective prescriptive programs for important economic decisions.

I. 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 at least two reasons. First, the problem is a hard one, even for an economist, so households might fail to compute the correct savings rate. Second, even if the correct savings rate were known, households might lack the self-control to reduce 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 middle-class 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 Steinmeier (1998), using the 1992 Health and Retirement Survey of households with heads of household 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 variety, however, in which self-control plays no role. Over the past decade, there has been a rapid change toward defined-contribution plans that require employees to actively join and select their own savings rate. For those workers who are eligible only for a defined-contribution 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 Steinmeier's analysis of workers who do not have pensions. The adequacy levels of their wealth and savings 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 offer only 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 TomorrowSM (or SMarT), and the basic idea is to give workers the option of committing themselves now to increasing their savings rate later, each time they get a raise. We report extensive data on one firm that implemented the program in 1998 and preliminary data on two other firms that implemented it recently.

We note that the null hypothesis predicted by the standard economic approach is that workers will have no interest in joining the SMarT 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.

II. A Prescriptive Approach to Increasing Savings Rates

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

¹ It is sometimes argued that this fact can be explained by selection effects, i.e., that 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 and important that workers sort themselves to firms primarily on this basis. Many other features of a job determine its attractiveness, and potential employees must make trade-offs. 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.

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 defined-benefit to defined-contribution savings plans is recent, there are as yet no 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 et al. (in press) report that two-thirds of their sample of 401(k) participants think that 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 to behavioral economics theorists (e.g., Laibson 1997; O’Donoghue 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 time consistent only 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} , since the difference in the value of the prizes exceeds the perceived costs of waiting. But as t approaches zero, the ratio of dis-

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 was computed to be a solution to the life cycle savings problem.

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

counted 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 naive. Sophisticated agents (modeled by Laibson) realize that they have hyperbolic preferences and take steps to deal with the problem, whereas naive 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 naivete 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 naive agents are, the more pronounced 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 savings 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 fixed-income 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 10-year panel of TIAA-CREF participants, finds a similar result. Nearly half of the participants made no changes to their plan over the 10-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 savings 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

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 because of the ease of trading via the Internet that was possible at both firms. But this increase in trading may also be attributable to rapidly rising stock prices during this period and the resulting excitement among individual investors.

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 costs of actively 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 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 proved 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 of over 90 percent (Choi et al., in press). But there is a downside 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 (see Choi et al. [in press], who explore these issues in depth). A goal of the SMarT plan is to obtain some of the advantages of automatic enrollment while avoiding some of the disadvantages.

On the basis of 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 are typically close to 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 et al. 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 pay increases may provide a propitious time to try to get workers to save more, since they are less likely to consider an increased contribution to the plan as a loss than they would at other times of the year.

In summary, 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 that they will get to it later. Our program to increase saving is aimed at these households.

III. The SMarT 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. On the basis of the principles discussed so far, we have proposed a program we call Save More TomorrowTM. The plan has four 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 dates 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

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

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 SMarT plan has many features that were included with the intention of making it attractive to employees who want to save more. It is not possible to say on theoretical grounds which features are most important, nor can theory tell us the ideal levels to select for many of the parameters that must be picked (e.g., the delay between the solicitation letter and the start of the program, the rate of increase, and the methods of soliciting and educating potential participants). Similarly, we cannot say a priori whether particular features, such as linking the increases in the savings rate to pay increases, are just one of many attractive components or are essential ingredients to success. We shall learn more about these questions over time as firms adopt the plan and provide data for analysis.

At this time we have three implementations on which we can report, each done rather differently. The particular design features were generally not selected by us but, rather, reflect the preferences of the firms that have adopted the plan. In this type of field research, we, the academic investigators, have quite limited control over many of the details, especially if compared with a laboratory environment. Nevertheless, it is not possible to study actual household savings behavior in a lab, so we are grateful for the data we are able to report here.

A. The First Implementation of SMarT: Midsize Manufacturing Company

The first implementation of the SMarT plan took place in 1998 at a midsize manufacturing company (which prefers to remain anonymous). The company, with the help of an investment consultant, selected the specific details of the implementation. Prior to the adoption of the SMarT 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 nondiscrimination 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.

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 the retirement savings plan. Of the 315 eligible

participants, all but 29 agreed to meet with the consultant and get his advice. On the basis of 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 constrain the program to increase the saving contribution by no more than 5 percent.⁷ 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 Internal Revenue Service (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. [Personal communication from Brian Tarbox, the investment consultant]

The participation data are reported in table 1. Of the 286 employees who talked to the investment consultant, only 79 (28 percent) were willing to accept his advice, even with the constraint that recommended increases were usually no more than five percentage points. For the rest of the participants, the planner offered a version of the SMarT plan as an alternative, proposing that they increase their saving rates by three percentage points each year, starting with the next pay increase. This was quite aggressive advice, since pay increases were barely more than this amount (approximately 3.25 percent for hourly employees and 3.50 percent for salaried employees). The pay increases were scheduled to occur roughly three months from the time the advice was being given.

⁷ Here and elsewhere, when we refer to a five-percentage-point increase, we are referring to an increase of percentage points, e.g., from a 2 percent saving rate to a 7 percent saving rate.

TABLE 1
PARTICIPATION DATA FOR THE FIRST IMPLEMENTATION OF
SMarT

Number of plan participants prior to the adoption of the SMarT plan	315
Number of plan participants who elected to receive a recommendation from the consultant	286
Number of plan participants who implemented the consultant's recommended saving rate	79
Number of plan participants who were offered the SMarT plan as an alternative	207
Number of plan participants who accepted the SMarT plan	162
Number of plan participants who opted out of the SMarT plan between the first and second pay raises	3
Number of plan participants who opted out of the SMarT plan between the second and third pay raises	23
Number of plan participants who opted out of the SMarT plan between the third and fourth pay raises	6
Overall participation rate prior to the advice	64%
Overall participation rate shortly after the advice	81%

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 SMarT 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 SMarT plan. More important, the majority of these participants did not change their mind once the savings increases took place. Only three participants (2 percent) dropped out of the plan prior to the second pay raise, with 23 more (14 percent) dropping out between the second and third pay raises and six more (4 percent) between the third and fourth pay raises.⁸ Hence, the vast majority of the participants (80 percent) have remained in the plan through four 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.

The impact of the SMarT plan on saving is shown in table 2.⁹ When

⁸ Interestingly, most of the employees who dropped out between the second and third pay raises worked for a single supervisor who apparently disapproved of the SMarT plan.

⁹ The data for each year refer only to those workers who are still employed by the company, so the sample shrinks over time from 315 to 229.

TABLE 2
AVERAGE SAVING RATES (%) FOR THE FIRST IMPLEMENTATION OF SMarT

	Participants Who Did Not Contact the Financial Consultant	Participants Who Accepted the Consultant's Recommended Saving Rate	Participants Who Joined the SMarT Plan	Participants Who Declined the SMarT Plan	All
Participants initially choosing each option*	29	79	162	45	315
Pre-advice	6.6	4.4	3.5	6.1	4.4
First pay raise	6.5	9.1	6.5	6.3	7.1
Second pay raise	6.8	8.9	9.4	6.2	8.6
Third pay raise	6.6	8.7	11.6	6.1	9.8
Fourth pay raise	6.2	8.8	13.6	5.9	10.6

* There is attrition from each group over time. The number of employees who remain by the time of the fourth pay raise is 229.

the investment consultant was brought into the company, the overall savings rate in the plan was 4.4 percent. The employees who did not want to talk to the consultant were saving more than the average, 6.6 percent. The group that accepted the advice of the consultant had been saving at exactly the overall company average, 4.4 percent, and after implementing the advice, they began saving 9.1 percent of their salary. At the end of our data collection period, that rate had slipped slightly to 8.8 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 three more raises, it was up 13.6 percent. Those participating in the SMarT plan ended up with a much higher saving rate than those who accepted the consultant's recommendation.

Of course, the implementation of the SMarT plan was not conducted as an experiment with random assignment to conditions. Participants selected themselves into the SMarT plan. In other circumstances, one might worry that the observed increase in savings rates might be attributable to some unmeasured "taste for saving" in the households that joined the SMarT plan; however, this worry seems unwarranted here on two counts. First, the SMarT 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 SMarT plan was offered only 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 SMarT plan.

The design of the study also rules out an information-based explanation for our results. Since the employees met with the investment consultant, they received useful information about proper savings rates, and this information quite possibly could affect their savings rates. All the employees who agreed to meet with the consultant received this information, however, 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 SMarT plan.

B. The Second Implementation of SMarT: Ispat Inland

The second implementation of the program took place in May 2001 at Ispat Inland, a large midwestern steel company. Ispat had heard about the SMarT plan and expressed to us an interest in increasing the saving rates among its 5,000 unionized employees. Ispat employees have received only one pay raise from the introduction of SMarT to date, so we can report only on the initial results at this time.

The implementation at Ispat was quite different from our first experience in that it was implemented with quite minimal resources. Most important, there was no financial consultant hired to meet one on one with employees. Instead, employees received a letter sent jointly by the human resources department and the union inviting them to join the SMarT program. There were no follow-up letters, no financial education seminars, and no other expenditures other than that single invitation letter and a few posters displayed in the cafeteria. In this implementation, the annual increase to the savings contribution rate was set at two percentage points every time they got a pay raise, with a cap on contribution rates set at 18 percent of salary. The first pay raise was scheduled for August 1, 2001, about two months after the solicitation letter was sent. The pay raise was to be 50 cents per hour, which amounted to roughly 2.5 percent of the average wage.

Participation in the program and the resulting saving rates are described in table 3. Even with this very inexpensive solicitation strategy, the program was popular with employees. Of the participants who were already enrolled in the 401(k) plan and were not already saving the maximum, 615 (18.1 percent) joined the SMarT plan. In addition, 165 employees joined SMarT who were not yet enrolled in the 401(k) plan; this was 8.2 percent of those employees who were eligible to participate in the 401(k) plan but had not yet enrolled. The lower take-up rate among the employees who were not currently in the 401(k) plan might be attributable to less interest in saving, but there is an additional con-

TABLE 3
AVERAGE SAVING RATES FOR ISPAT INLAND (%)

	EMPLOYEES WHO WERE ALREADY SAVING ON MAY 31, 2001		EMPLOYEES WHO WERE NOT SAVING ON MAY 31, 2001		ALL ELIGIBLE EMPLOYEES (<i>N</i> =5,817)
	Joined SMarT (<i>N</i> =615)	Did Not Join SMarT (<i>N</i> =3,197)	Joined SMarT (<i>N</i> =165)	Did Not Join SMarT (<i>N</i> =1,840)	
Pre-SMarT (May 2001)	7.62	8.62	.00	.00	5.54
First pay raise (October 2001)	9.38	8.54	2.28	.26	5.83

NOTE.—The sample includes 5,817 employees who are eligible to participate in the 401(k) plan and have remained with the company from May 2001 through October 2001. The sample includes 414 employees who were already saving at the maximum rate of 18 percent, although they were not allowed to join the SMarT program. The reported saving rates represent the equally weighted average of the individual saving rates.

tributing factor. Those who were not in the plan might have ignored the letter altogether. The letter came with the heading “important information about your 401(k) account,” a teaser that would not be of particular interest to employees who were not in the plan.

The immediate effect on savings was about what might be expected. Those joining SMarT increased their saving rates by roughly 2 percent, whereas those not joining the program did not change their saving rates much. If the experience in the first implementation is repeated here and few employees drop out of the SMarT plan, then saving rates will continue to increase whenever the employees get raises.

C. *The Third Implementation of SMarT: Philips Electronics*

The third implementation of SMarT took place at two divisions (Divisions A and O) of Philips Electronics in January 2002, with the first saving increase taking place on April 1, 2002.¹⁰ The remaining 28 divisions of Philips served as a control. Invitation letters were sent to 815 “non–highly compensated” employees whose saving rates were below 10 percent. Everything was done the same way at the two divisions except for the following: Employees at Division A were given the option of attending educational seminars devoted to retirement savings (including a description of the new SMarT plan) but were not offered any one-on-one meetings. For the employees at Division O, attendance at the financial education seminar was strongly encouraged. The seminar was described to the employees as “required,” although there was no penalty for failing to meet the requirement. Whether because of the “require-

¹⁰ Additional details on the implementation at Philips are available at http://institutional.vanguard.com/pdf/SMarT_112002.pdf.

ment” or other reasons, 60 percent of the employees attended the seminar, whereas only 40 percent did so in Division A. The employees in Division O were also offered the opportunity to have a one-on-one meeting with a certified financial planner. The average saving rates prior to SMarT were quite similar at the two divisions, 3.12 percent and 3.74 percent for Divisions A and O, respectively, both rates slightly higher than the saving rates in the rest of the Philips divisions (2.90 percent). But the two divisions are different along many other dimensions, making direct comparisons difficult. For example, Division A is in the technology business, is located in the desert Southwest, and was suffering through a severe recession at the time of the implementation, whereas Division O focuses on consumer products, is located in the Pacific Northwest, and has been doing well economically.¹¹ Thus the two divisions do not represent a true controlled experiment in comparison with each other, though they can reasonably be compared with the other control divisions, at least in terms of saving rates.

There were two notable differences between the implementation at Philips and the previous two trials described above. First, increases in savings were not necessarily linked to pay raises. Instead, employees were told that if they joined the plan, their saving rates would go up on April 1 of each year whether or not they received a pay raise. Pay raises do tend to occur on April 1, but the employees could not be sure that the extra contribution to the savings plan would come out of their raise. Second, employees were allowed to pick the rate at which their savings would increase: one, two, or three percentage points per year. Those who joined the plan but did not choose a rate of increase were assigned a 2 percent rate of increase.¹² Fifty-four percent of the SMarT enrollees elected an annual increase of 1 percent, 35 percent elected the default of 2 percent, and the remaining 11 percent elected 3 percent. Regardless of the chosen annual increase, the annual increases will stop once the participant reaches a saving rate of 10 percent.

The resulting saving rates are displayed in table 4. As expected, not much is happening at the remaining 28 divisions of Philips Electronics that served as our control group. In contrast, saving rates for those who were already enrolled in the 401(k) plan and joined the SMarT plan went up, as expected, by about 1.5 percent (the weighted-average pro-

¹¹ Division A is now closing down, so the long-term results of the SMarT plan will not be available.

¹² There are pros and cons to offering this choice to participants, as opposed to just picking a single rate of increase. The obvious advantage is that employees can select the rate of increase they like best. The disadvantage is that simply being forced to make such a choice adds another layer of complexity that could discourage some potential enrollees. We included the default 2 percent rate of increase with the goal of mitigating this potential impediment to enrolling. Only a controlled experiment will be able to determine whether the pros of offering choice outweigh the cons.

TABLE 4
AVERAGE SAVING RATES (%) FOR PHILIPS ELECTRONICS

DATE	EMPLOYEES WHO WERE ALREADY SAVING IN DECEMBER 2001		EMPLOYEES WHO WERE NOT SAVING IN DECEMBER 2001		ALL EMPLOYEES
	Joined SMarT	Did Not Join SMarT	Joined SMarT	Did Not Join SMarT	
	A. Control Group				
Observations	7,405		7,053		14,458
Pre-SMarT (December 2001)	5.65		.00		2.90
Post-SMarT (March 2002)	5.76		.70		3.29
B. Test Group (Divisions A and O Combined)					
Observations	180	339	36	260	815
Pre-SMarT (December 2001)	5.26	5.38	.00	.00	3.40
Post-SMarT (March 2002)	6.83	5.72	5.03	1.55	4.61
C. Division A					
Observations	66	190	10	163	449
Pre-SMarT (December 2001)	5.47	5.48	.00	.00	3.12
Post-SMarT (March 2002)	7.32	5.97	6.80	1.54	4.38
D. Division O					
Observations	114	149	26	77	366
Pre-SMarT (December 2001)	5.14	5.25	.00	.00	3.74
Post-SMarT (March 2002)	6.55	5.41	4.35	1.58	4.89

NOTE.—The “test” group consists of individuals at Divisions A and O.

grammed increase among those already saving). The savings rate went up more dramatically for those employees who simultaneously enrolled in the 401(k) and SMarT. Interestingly, there also seemed to be a spillover effect on those who did not join the SMarT plan. In the two experimental divisions in which SMarT was introduced, even those employees who did not join SMarT increased their savings rates more than was observed in the control group.¹³

At the time of this writing (summer 2003), the second raise has occurred in Division O, and we have some preliminary data on attrition rates from the SMarT program. Of those who originally joined the program, 13.5 percent have left Philips because they either quit or were terminated. Of those remaining in the plan, eight employees (5.4 percent of the original participants, 6.2 percent of those still working at Philips) dropped out before the second raise, but another five em-

¹³ This pattern is consistent with evidence by Duflo and Saez (2000) on peer effects.

ployees joined the plan. This experience of low dropout rates is comparable to that in the first implementation and suggests that, over time, savings rates will continue to rise.

In this implementation, we were given access to some demographic information about the employees as well as information about how the plan was administered in each division. The participation rates in SMarT were quite different in the two divisions, as shown in table 5. In Division A, only 16.9 percent of the division's employees joined the program (76/449), whereas Division O had a take-up rate of 38.3 percent (140/366). One potential explanation for this difference is that the employees at Division O had the opportunity to meet with a certified financial planner. In fact, 41.8 percent of the employees at Division O met with the financial planner, and 81 percent of those who attended such a meeting actually joined SMarT. Of course, electing to meet with the planner might by itself signal a desire to save, so it is not possible to ascertain the incremental effect of the financial planner on either saving rates or SMarT participation.

Table 5 also provides some basic information on who joins the SMarT plan. Neither gender nor age appears to be an important determining factor. Employees with four to five years of tenure working for Philips were the most likely to join, as were those with annual incomes of less than \$50,000.

At this stage, there are some preliminary lessons that can be drawn from the Philips experience. First, the SMarT design feature linking savings increases to pay increases, while desirable, may not be essential. This is important, since some firms find it difficult to coordinate the savings plan with the salary increases. Second, one-on-one meetings with a financial planner appear to be a very effective (though costly) recruitment tool, though selection problems make it difficult to parse out the precise value of this intervention.

IV. SMarT and Savings Adequacy

Obviously, the SMarT plan can produce dramatic increases in saving rates. In the first implementation, those who joined the plan more than tripled their saving rates in 28 months. This raises the question of what effect SMarT 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 using lessons from our first implementation. We focus on the first implementation since it has the longest track record.

We do not have demographic or financial information about the em-

TABLE 5
PARTICIPATION RATES IN THE SMarT PROGRAM AT PHILIPS ELECTRONICS

Explanatory Variable	Total Number of Employees	Number of Employees in the Test Group	Number of Employees Joining SMarT	SMarT Participation Rate (%)
Entire sample	15,273	815	216	26.5
Sex:				
Missing	9,355	430	92	21.4
Female	2,191	146	51	34.9
Male	3,727	239	73	30.5
Age:				
Missing	3,598	134	33	24.6
20s	1,880	162	46	28.4
30s	3,990	264	66	25.0
40s	3,944	178	49	27.5
50s	1,861	77	22	28.6
Tenure (in years):				
0–1	1,953	103	17	16.5
2–3	3,096	200	58	29.0
4–5	2,064	137	49	35.8
6–10	3,087	262	69	26.3
11+	5,073	113	23	20.4
Income:				
Missing	4,207	10	0	.0
<\$25,000	1,786	155	56	36.1
\$25,000–50,000	4,296	362	106	29.3
\$50,000–75,000	2,386	134	27	20.1
\$75,000+	678	53	11	20.8
Saving rate (prior to SMarT):				
0%	7,351	296	36	12.2
1–5%	1,914	162	62	38.2
6%	4,931	304	101	33.2
7–9%	1,079	53	17	32.1
Division:				
A	449	449	76	16.9
Control	14,458	0	0	
O	366	366	140	38.3
Participated in education seminar:				
No	389	389	20	5.6
Yes	426	426	196	46.0
Met with financial advisor:				
No	213	213	16	7.5
Yes	153	153	124	81.0
Registered Web user:				
No	12,161	663	162	24.4
Yes	3,112	152	54	35.5

NOTE.—The initial sample included 46,873 individual-year observations (excluding highly compensated employees). We first required that all the individuals be present before and after the implementation of the SMarT program, which reduced the number to 20,122 individuals. Next, we eliminated those who switched between the test and control groups, leaving us with 20,103 individuals. We also eliminated those saving more than 10 percent of their pay because they were not allowed to join SMarT, resulting in 15,274 individuals. Of the remaining 15,274 individuals, most are in the “control” group, and they were not offered the SMarT program. The “test” group consists of individuals at the A and O Divisions.

employees in our study, so we need to make some assumptions about their household financial situations in order to calculate the likely effects of joining the SMarT plan. We make calculations for hypothetical 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 for some of the larger 401(k) plans they administer. In particular, we calculate the account balances of people of a similar age, income, and savings rate. 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 that employees are saving 4 percent in the 401(k) plan when they join the SMarT plan and that saving rates are capped at 14 percent. We also assume that the employer matches the employee's 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, on the basis of data from John Hancock Financial Services (the Sixth Defined Contribution Plan Survey [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. Finally, we assumed the statutory benefits from social security. We then use software provided by Financial Engines to estimate the distribution of retirement income that can be expected on the basis of these assumptions. The Financial Engines software provides several points on a probability distribution of retirement income. We use the fiftieth percentile of this distribution to compute expected income replacement rates, that is, the ratio of retirement income to preretirement income.

Table 6 reports retirement income replacement rates for various age and income combinations. The issue of savings adequacy is well studied by economists, but there is no agreement on a single number as the appropriate replacement rate (see Boskin and Shoven [1987], Bernheim [1993], and Gustman and Steinmeier [1998] for discussions of this issue). Still, most economists writing on this issue consider replacement rates near 100 percent adequate and judge replacement rates below 70 percent to be too low.¹⁴

Panel A of table 6 shows the expected income replacement rates for our employees before they join the SMarT plan, all of which are between

¹⁴ One might think that a 100 percent replacement rate would be too high, suggesting that agents are very patient. However, survey evidence suggests that households desire an increasing consumption profile. Laibson (1999) offers a cogent discussion of this issue and 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.

TABLE 6
MEDIAN INCOME REPLACEMENT RATIOS (%)

INCOME	AGE			
	25	35	45	55
A. Pre-SMarT				
\$25,000	57	57	56	55
\$50,000	51	51	51	54
\$75,000	48	49	46	43
B. Post-SMarT				
\$25,000	108	90	75	63
\$50,000	98	83	70	62
\$75,000	90	77	63	50

NOTE.—The 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 4 percent before SMarT and 14 percent thereafter, employer match of 50 cents on the dollar up to 6 percent, portfolio mix of 60 percent stocks and 40 percent bonds, and retirement age of 65.

43 and 57 percent. 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 SMarT plan, especially for those joining the plan when young. Obviously, increasing the savings rate is less effective when one starts at 55 than at 25. Still, expected replacement rates exceed 100 percent in just one cell (108 percent replacement for those making \$25,000 per year who join the plan at age 25), so there does not appear to be evidence that we have induced people to save too much. Furthermore, if the stock market returns are exceptionally high, workers nearing retirement can always reduce savings rates or plan an earlier retirement if they have higher retirement benefits than they expected.

V. The Potential Effect of SMarT on the U.S. Personal Savings Rate

The U.S. personal savings rate is currently close to zero. Some macroeconomists consider this rate too low and have advocated government intervention to increase the savings rate. We do not take any stand on whether such policies are good for the economy but, instead, ask a different question. If it were desirable to increase the personal savings rate, could widespread adoption of the SMarT plan make a substantial contribution to meeting this goal?

To determine the potential impact of widespread adoption of SMarT, we begin by characterizing how much employees are saving now in their 401(k) plans. To do so, we utilize a data set from Hewitt Associates that includes demographic and account balance information on the partic-

ipants in 15 large companies, covering a total of 539,516 employees. On the basis of comparisons with data from Fidelity (a 2001 report on corporate defined-contribution plans) and John Hancock Financial Services (the 1999 Sixth Defined Contribution Plan Survey), two other large 401(k) service providers, we believe that our sample is representative of employees who work for large companies. Consequently, we think that it can serve as a basis for some rough estimates on the potential contribution SMarT can make to increasing employee savings rates.

Starting from the baseline behavior we observe now, we make calculations of changes in savings rates over a 10-year period for various implementation strategies. Specifically, we consider three hypothetical implementation strategies, each matched with increases in the savings rate of 1, 2, or 3 percent per year, giving us nine configurations to examine. We start each plan at a 5 percent savings rate, approximately the average in the Hewitt data. We then simulate the impact of adding specific implementations of SMarT. In all the simulations, we assume that 5 percent of enrollees drop out of SMarT each year, leaving their savings rate at the level they had obtained up to that point.

The first two implementation strategies we consider are based on the experiences we have had in the implementations described above. Plan A is based on the first implementation, which used one-on-one interactions with a financial consultant. On the basis of the results in that company, we assume that 80 percent of those who are currently participating in the savings program will join the SMarT plan, and half of those who are not enrolled will join. Plan B is based on the experience at Ispat, where the SMarT plan was marketed to employees only with a single direct-mail campaign rather than personal contact. This approach is much less costly but is also less effective in reaching potential enrollees. In this scenario, 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 SMarT program with automatic enrollment. Specifically, we assume that all employees would be enrolled in the SMarT plan unless they opted out. Those who are not currently participating in the 401(k) plan would be enrolled, and their initial saving rate would be the savings incremental rate (i.e., 1, 2, or 3 percent). On the basis of our experience and that of Madrian and Shea (1999) and Choi et al. (in press), for plan C we estimate that 90 percent of the employees would join the program in this design (i.e., only 10 percent would opt out). 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.¹⁵

¹⁵ For simplicity, our calculations exclude the effects of employer contributions and

TABLE 7
PROJECTED SAVING RATES (%)

SMarT Annual Increments (%)	PROJECTED SAVING RATES WITH SMART IN YEAR (%)						
	0	1	2	3	4	5	10
A. One-on-One Interaction with a Financial Consultant							
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
B. One-Shot Mailing							
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
C. Automatic Enrollment							
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

The results of our projections are displayed in table 7. 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.¹⁶ This means that there is considerable opportunity for the SMarT program to increase the saving rate. With plan A, which uses one-on-one interaction with a financial consultant and the 2 percent per year rate of increase, the SMarT program could boost the overall saving rate from 5.0 percent to 9.7 percent within five years (see panel A). When one switches 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

employee turnover. These omissions create biases in opposite directions. On one hand, including employer contributions would increase the estimated effect of the SMarT 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 SMarT program unless the employee moves to another firm with the SMarT plan in effect.

¹⁶ The IRS limit for the year 2000 was the lower of \$10,500 or 25 percent of income.

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 (data from the National Income and Product Accounts: Personal Income and Its Disposition: <http://www.bea.doc.gov/briefrm/tables/ebr6.htm>). Since the current personal savings rate is hovering near zero, this is a substantial increase. Furthermore, in contrast to other approaches to increasing the employee savings rate, such as increasing the maximum allowable contribution, much of the gains from the SMarT 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.

VI. Conclusions

The initial experience with the SMarT plan has been quite successful. Many of the people who were offered the plan elected to use it, and a majority of the people who joined the SMarT plan stuck with it. Consequently, in the first implementation, for which we have data for four annual raises, SMarT participants almost quadrupled their saving rates. Of course, one reason why the SMarT plan works so well is that inertia is so powerful. Once people enroll in the plan, few opt out. The SMarT 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 financial consultant involved in the first implementation has noted, in hindsight it would have been better to offer the SMarT plan to all participants, even those who were willing to make their initial savings increase more than the first step of the SMarT plan. Very few of these eager savers ever got around to changing their savings allocations again, whereas the SMarT plan participants were already saving more than they were after just 16 months (see table 2)

Some economists have criticized practices such as automatic enrollment and the SMarT plan on the grounds that they are paternalistic, a term that is not meant to be complimentary. We agree that these plans are paternalistic, but since no coercion is involved, they constitute what Sunstein and Thaler (2003) call “libertarian paternalism.”¹⁷ Libertarian paternalism is a philosophy that advocates designing institutions that help people make better decisions but do not impinge on their freedom to choose. Automatic enrollment is a good example of libertarian paternalism. Notice that firms must decide what happens to employees who take no action with respect to joining the savings plan. Traditionally, employees who did nothing were presumed *not* to want to join the plan.

¹⁷ For a brief summary of this idea, see Thaler and Sunstein (2003).

Automatic enrollment simply changes that presumption. Neither arrangement infringes on choice (so both are libertarian), but one produces higher savings rates and so might be considered paternalistic. The SMarT plan is even less intrusive than automatic enrollment since participants have to take some action to enroll, and it is even more successful at getting people to save. So, we plead guilty to the charge of trying to be paternalistic, but since we are striving for libertarian paternalism, we do not think that it should be considered objectionable.

Finally, we hope that 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 a 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, George. 1975. "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control." *Psychological Bull.* 82 (4): 463–96.
- Ameriks, John, and Stephen Zeldes. 2000. "How Do Household Portfolio Shares Vary with Age?" Working paper. New York: Columbia Univ.
- Bernheim, B. Douglas. 1993. *Is the Baby Boom Generation Preparing Adequately for Retirement?* Plainsboro, N.J.: Merrill Lynch.
- Bernheim B. Douglas, Daniel M. Garrett, and Dean M. Maki. 1997. "Education and Saving: The Long-Term Effects of High School Financial Curriculum Mandates." Working Paper no. 6085 (July). Cambridge, Mass.: NBER.
- Boskin, Michael J., and John B. Shoven. 1987. "Concepts and Measures of Earnings Replacement during Retirement." In *Issues in Pension Economics*, edited by Zvi Bodie, John B. Shoven, and David A. Wise. Chicago: Univ. Chicago Press (for NBER).
- Choi, James J., David Laibson, Brigitte Madrian, and Andrew Metrick. In press. "For Better or for Worse: Default Effects and 401(k) Savings Behavior." In *Perspectives on the Economics of Aging*, edited by David A. Wise. Chicago: Univ. Chicago Press (for NBER).
- Choi, James J., David Laibson, and Andrew Metrick. 2000. "Does the Internet Increase Trading? Evidence from Investor Behavior in 401(k) Plans." Working Paper no. 7878 (September). Cambridge, Mass.: NBER.
- Duflo, Esther, and Emmanuel Saez. 2000. "Participation and Investment Decisions in a Retirement Plan: The Influence of Colleagues' Choices." Working paper. Cambridge: Massachusetts Inst. Tech.
- Farkus, Steve, and Jean Johnson. 1997. *Miles to Go: A Status Report on Americans' Plans for Retirement*. New York: Public Agenda.
- Gustman, Alan L., and Thomas L. Steinmeier. 1998. "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study." Working Paper no. 6681 (August). Cambridge, Mass.: NBER.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *A.E.R.* 76 (September): 728–41.

- . 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *J.P.E.* 98 (December): 1325–48.
- Laibson, David I. 1997. "Golden Eggs and Hyperbolic Discounting." *Q.J.E.* 112 (May): 443–77.
- . 1999. Discussion of "The Adequacy of Household Saving" by Eric Engen, William Gale, and Cori Uccello. *Brookings Papers Econ. Activity*, no. 2, pp. 174–77.
- Loewenstein, George, and Jon Elster, eds. 1992. *Choice over Time*. New York: Sage.
- Madrian, Brigitte C., and Dennis Shea. 1999. "The Power of Suggestion: An Analysis of 401(k) Participation and Saving Behavior." Working paper. Chicago: Univ. Chicago, Grad. School Bus.
- O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *A.E.R.* 89 (March): 103–24.
- . 2001. "Choice and Procrastination." *Q.J.E.* 116 (February): 121–60.
- Raiffa, Howard. 1982. *The Art and Science of Negotiation*. Cambridge Mass.: Harvard Univ. Press.
- Samuelson, William, and Richard J. Zeckhauser. 1988. "Status Quo Bias in Decision Making." *J. Risk and Uncertainty* 1 (March): 7–59.
- Shafir, Eldar, Peter Diamond, and Amos Tversky. 1997. "Money Illusion." *Q.J.E.* 112 (May): 341–74.
- Shefrin, Hersh M., and Richard H. Thaler. 1988. "The Behavioral Life-Cycle Hypothesis." *Econ. Inquiry* 26 (October): 609–43.
- Sunstein, Cass R., and Richard H. Thaler. 2003. "Libertarian Paternalism Is Not an Oxymoron." *Univ. Chicago Law Rev.* 70 (Fall): 1159–99.
- Strotz, Robert H. 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *Rev. Econ. Studies* 23 (3): 165–80.
- Thaler, Richard H. 1981. "Some Empirical Evidence on Dynamic Inconsistency." *Econ. Letters* 8 (3): 201–7.
- Thaler, Richard H., and Hersh M. Shefrin. 1981. "An Economic Theory of Self-Control." *J.P.E.* 89 (April): 392–406.
- Thaler, Richard H., and Cass R. Sunstein. 2003. "Libertarian Paternalism." *A.E.R. Papers and Proc.* 93 (May): 175–79.
- Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *J. Risk and Uncertainty* 5 (October): 297–323.