Intertemporal Choices for Health

Justin S. White, Stanford University
William H. Dow, UC Berkeley

July 1, 2014

This chapter was prepared for a book titled *Behavioral Economics and Public Health*, edited by Christina A. Roberto and Ichiro Kawachi, to be published by Oxford University Press.

INTRODUCTION TO INTERTEMPORAL CHOICE

It is New Year's Eve. In keeping with tradition, many revelers resolve to lose weight, to find time for exercise, or to kick some unpleasant habit. When we encounter these individuals many months later, we find their lives and their health behaviors unchanged. Why is it so difficult for individuals to follow through on their plans and goals? What might we as social scientists and public health researchers conclude about the preferences and choices of these individuals? This chapter addresses the wedge between the health-related plans and goals of individuals and their subsequent actions. In so doing, we also consider the wedge between how economists have traditionally understood issues of intertemporal choice and how newer cohorts of economists have applied the principles of behavioral economics to understand the topic. In comparing standard economic theory and behavioral economic theory, we highlight the unique contributions of behavioral economic insights to a fuller conception of intertemporal choices for health.

By intertemporal choice, we refer to any decision that impacts a person's future welfare. The consumption of many health-related goods does not take place in a single instance, but over time. Take smoking as an example. A typical life-course trajectory might be the following. At a young age, Maria decides to take a first puff. Soon, she sneaks a cigarette after school every day. In five years, she has developed a habit of smoking half a pack each day. Many health behaviors such as smoking involve the accretion of a series of incremental decisions. Each day, Maria wakes up and chooses whether or not to smoke, in essence whether to sustain the habit or to try to kick it. Economists theorize that in making these daily decisions Maria is solving a series of intertemporal choice problems in which she trades off the current pleasures and harms of smoking against the future pleasures and harms of having smoked. The difference between how traditional economics and behavioral economics view Maria's smoking behavior boils down to certain assumptions about how Maria arrived at a decision to smoke. We consider those assumptions in detail below.

The field of behavioral economics arose to help understand anomalies that are difficult to explain using the framework of standard (i.e., "neoclassical") economic theory. Among the most prominent anomalies are those pertaining to intertemporal choices, and many of the most glaring examples relate to health. The Nobel-winning economist Thomas Schelling (1980) wrote, "People behave sometimes as if they had two selves, one who wants clean lungs and long life and another who adores tobacco, or one who wants a lean body and another who wants dessert." In the intervening decades, we have learned a great deal about how this time-inconsistent behavior pervades an assortment of health-related activities. Behavioral economic concepts have been applied in health most frequently to the consumption of habit-forming goods such as alcohol and unhealthy foods, although the power of these concepts extends far beyond this subset. For example, non-adherence to medications can often be attributed to the intertemporal choices of patients: an inability or unwillingness to weather the short-term side effects of medications, a miscalculation about the harms caused by non-adherence, or a limited ability to remember when or how much medication to consume. The implications for public health practitioners of not addressing problems of intertemporal choice are significant. For example, non-adherence may have serious consequences for personal health and the health system as a whole, accounting for one-third to two-thirds of all medication-related hospital admissions in the U.S. (Osterberg and Blaschke, 2005).

To take another example, more than one in three patients is not screened for colorectal cancer as recommended by federal guidelines (Klabunde et al., 2012). This situation shares some of the same characteristics as medication non-adherence. Individuals may not wish to undergo an invasive and unpleasant colonoscopy; they may not fully realize the value of the screening; or they may not remember when it is time for a screening exam or how often to get one. Elucidating the existence and importance of each of these pathways for health-related activities is an important challenge for researchers.

Overview of the chapter

Our focus in this chapter is on the systematic—that is, potentially foreseeable—ways in which individuals' intertemporal health choices depart from the predictions of standard economic theory. By "standard" theory we refer to a model (elaborated in the next section) of the rational agent from neoclassical economics who is perfectly informed, forward-looking and consistent in his preferences. Section 2 begins with a description of the standard economic model of rational choice. We then describe three classes of deviations from the standard model, relying on a framework developed by DellaVigna (2009): (1) nonstandard preferences, (2) nonstandard beliefs, and (3) nonstandard decision-making. We often adopt the first person in our illustrations in recognition that we are as susceptible to cognitive mistakes as anyone.

The first and most studied class is the apparent display of preferences that are inconsistent with the standard model. We focus on time preferences that generate self-control problems. Situations that result in immediate pleasure or pain often lead to self-control problems, as discussed in Section 3. A desire to avoid immediate pain may lead us to delay quitting smoking or to avoid our annual flu shot. Likewise, a desire to indulge in immediate pleasures may lead us to overeat or to engage in drug use or unprotected sex.

The second class of deviations involves nonstandard beliefs that individuals hold about the future. In Section 4, we consider two types of mispredictions that impact health decisions. First, when in a state of heightened emotion, arousal, or hunger, we fail to predict how we will feel in the future after we have emerged from the state. While hungry, we fail to consider how many entrées we would have ordered in a more sated state. In the heat of the moment, we fail to consider fully our underlying preferences for using contraception. Second, we tend to be overly confident about our abilities. For example, we tend to view ourselves as impervious to illness and perfectly capable of resisting temptations in the future.

The third class of deviations from the standard model involves nonstandard decision-making. In Section 5, we tackle the cognitive limitations of human attention and memory and the ways that limited attention and memory can undermine our health behavior.

After introducing and illustrating the ways that our preferences, beliefs, and decision-making may systematically diverge from the standard model, in Section 6 we turn to the implications for policy and the design of decision aids. What tools can be leveraged to help individuals deal with complex and challenging choice environments? First we discuss one of the most celebrated approaches to promoting time-consistent choices: a strategy of precommitment. Next, we consider more broadly how public health programs, and in particular the incentives within programs, can be designed to combat cognitive biases, or decision errors that fail to promote the person's well-being. Then, we discuss the special role played by taxation as a counterweight to the urges to partake in undesirable health behaviors. Finally, we concentrate on vulnerable populations to discuss how cognitive biases can perpetuate disparities across socioeconomic strata and lead to life-long consequences when initiated in adolescence and young adulthood.

This chapter considers several theoretical constructs from behavioral economics that affect intertemporal choices for health. Yet, there are many more that we do not have space to explore, such as heuristical thinking, social preferences, and reference dependence. These topics are considered in other chapters in this volume. Daniel Kahneman and Amos Tversky's (2000) *Choices, Values, and Frames* is a good starting point for a general discussion of these topics. We have also tried to be somewhat selective in our presentation of the literature, as it has grown in recent years. We direct most of our attention to evidence drawn from field settings, rather than laboratory settings. By doing so, we hope to make clear the applicability

of this line of inquiry for real-world public health research and practice and the expansive implications of behavioral economics for intertemporal choices across the entire spectrum of health behavior.

RATIONAL CHOICE THEORY

Rational choice underpins much of microeconomic theory. In this section, we provide a brief overview of the standard economic theory of rational choice and its application to public health. Then, we describe certain critiques of the theory that have led in part to the emergence of alternative models of behavior.

Description of the standard model

Individuals are considered to be rational if they make choices so as to obtain the most utility (satisfaction) from a given set of resources. Satisfaction or utility depends on tastes or preferences for different goods.

Individuals are assumed to have well-defined preferences. This means that for any pair of goods, say fish and steak, a person can always say that fish is preferred to steak, steak is preferred to fish, or the two are equally attractive. These preferences are often assumed to be stable over time, because this assumption simplifies the researcher's job.

Individuals trade off present utility for future utility using a weighting function called the discount factor. The same resource is more valuable in the present than in the future, and the discount factor represents the degree to which a person down-weights future consumption relative to present consumption. A peculiar feature of the standard model is that the discount function is typically assumed to be constant, such that short-run tradeoffs are treated the same as long-term tradeoffs. This implies that individuals have time-consistent preferences. If a rational consumer prefers to visit the gym tomorrow but not today, then all else equal he will always go to the gym when tomorrow comes. Time-inconsistency, or preference reversals, in which a person sleeps in tomorrow are not usually possible in the standard model.[1]

The standard model also makes assumptions about the beliefs of individuals. They are forward-looking individuals who plan ahead, with full information about available options. Moreover, individuals hold rational expectations about their own future behavior. This means that individuals' predictions of the likelihood of future outcomes and the payoffs that future outcomes yield are correct on average.

---

[1] Paul Samuelson (1937) recognized the restrictiveness of a constant discount function when he introduced discounted utility, although exponential discounting nevertheless became standard economic practice.

Decision-making in the standard model involves the maximization of lifetime expected utility. A person uses all available information, factoring in the uncertainty surrounding all possible events, to choose the consumption plan that will produce the maximum level of utility over the person's lifetime. A person is assumed to behave as if he follows this principle, even if he does not explicitly perform such calculations in his head. If a new alternative presents itself, the person integrates it into his current consumption plan before making a choice.

A simple utility model

To show how economists apply the assumptions of the standard model in practice, we present a simple model of utility. Suppose that a person lives for T periods, $t = 0, 1, 2, \dots, T$. For simplicity, you may think of each period as one day. In each period t, a person receives instantaneous utility as a function of the goods consumed during that period. Imagine a person only derives utility from a single good x. The person's instantaneous utility at time t is denoted $u(x_t)$, which we shorten to u_t for simplicity. Thus, in periods 1, 2, and T, a person's instantaneous utility would be u_1, u_2, and u_T.

Building on this framework, we can determine a person's lifetime utility U as a function of consumption in each time period $U = U(x_0, x_1, \dots, x_T)$. The person's lifetime utility function is the sum of utility in period 1 and discounted future utility in periods 2 through T.

$$\begin{aligned} U &= u_0 + \delta u_1 + \delta^2 u_2 + \delta^3 u_3 + \cdots + \delta^T u_T \\ &= \sum_{t=0}^{T} \delta^t u_t \end{aligned} \tag{1}$$

The most commonly used discount function is the exponential function δ^t, as used here. The discount factor δ represents the degree to which a person discounts future consumption relative to current consumption. For example, owning a bicycle today is more valuable than owning one tomorrow by a factor δ; owning the bicycle tomorrow is more valuable than the day after tomorrow by a factor δ; and owning one today versus the day after tomorrow by a factor $\delta^2 = \delta \times \delta$.

A key property of *exponential discounting* is that the discount factor from any period t to the subsequent period $t + 1$ is constant.[2] A person will discount the same between today and tomorrow as between tomorrow and the day after tomorrow. An implication of this constant discount factor is that the choices of an "exponential discounter" are always time-consistent.

[2] This property is called a constant marginal rate of substitution.

A standard model incorporating uncertainty (casual readers may skip this section!)

We now add one more layer of complexity to our utility function in order to capture the notion that choices have uncertain outcomes. Individuals make choices based not only on the outcomes of a choice but also on the probability of that outcome being realized. For example, Anya is a 40-year-old woman trying to decide whether to get a mammogram for breast cancer screening. She must consider the benefit from a mammogram in the presence and absence of cancer. She also needs to consider her probability of having breast cancer.

Suppose that Anya has a 25% probability of being diagnosed with breast cancer. Further assume that if she gets a mammogram her instantaneous utility is 2 if diagnosed with cancer and 10 if not diagnosed with cancer. Anya's expected utility from the mammogram would be: $(0.25 \times 2) + (0.75 \times 10) = 8$. Assume that her expected utility if she does not get a mammogram is 10. In this case, Anya would choose not to get a mammogram, as it yields greater utility.

Formally, let $p(s)$ be the probability of outcome s, and let $u(x|s)$ be the person's instantaneous utility from consuming x given outcomes $s = 1, 2, \dots, S$. In this simple case, where we have omitted any time component, expected utility $\mathrm{E}(U)$ is the product of the payoff from each outcome and the probability of each outcome, just as for Anya above:

$$\begin{aligned} E(U) &= p(1)\, u(x|1) + p(2)\, u(x|2) + \cdots + p(S)\, u(x|S) \\ &= \sum_{s=1}^{S} p(s) u(x|s) \end{aligned} \tag{2}$$

Under rational expectations, a person's beliefs about the probability of each outcome $p(s)$ are accurate. In other words, Anya is fully informed about her cancer risk. If we extend the model to capture future mammography choices as well (as in Equation (3) below), Anya is also assumed to accurately predict her future cancer risk.

We combine the concepts from Equations (1) and (2) to show a simple model of how a person makes intertemporal choices according to the standard model. A person who is about to enter period 0 faces uncertain choices in each time period $t = 0, 1, \dots, T$ that are subject to probabilities $p(s_t)$ for outcomes s_t. This person maximizes lifetime expected utility by maximizing expected utility in each period.

$$\max_{x_t} \underbrace{\sum_{s_0} p(s_0)\, u(x_0|s_0)}_{\substack{\text{Expected utility} \\ \text{in period 0}}} + \underbrace{\delta \sum_{s_1} p(s_1)\, u(x_1|s_1)}_{\substack{\text{Discounted expected} \\ \text{utility in period 1}}} + \cdots + \underbrace{\delta^T \sum_{s_T} p(s_T)\, u(x_T|s_T)}_{\substack{\text{Discounted expected} \\ \text{utility in period } T}}$$

$$= \max_{x_t} \sum_{t=0}^{T} \delta^t \sum_{s_t} p(s_t)\, u(x_t|s_t) \tag{3}$$

Equation 3 states that a person maximizes the expected utility from her current and future consumption of good x, where future utility is discounted by discount factor δ.

Criticisms of the standard model

Standard utility theory has been criticized along several lines. Some researchers have pointed to the long list of apparently self-defeating behaviors that are at odds with the predictions of the theory. The sheer prevalence of drug addiction, gambling, unsafe sexual behavior, obesity, and many others belies the existence of utility-maximizing behavior. Individuals who struggle with these conditions often feel regret and shame, and often invest large sums of money and time trying to change their behavior as evidenced by the profitability of the weight loss and self-help industries and the popularity of programs such as Alcoholics Anonymous. Throughout this chapter, we highlight evidence that does not conform to the predictions of the standard model.

Proponents of the standard model have tried to marshal evidence in support of the standard model. A notable case is the application of rational choice theory to addictive behavior. In the "rational addiction model," individuals choose to become addicted and to maintain an addiction after determining that it enhances their utility (Becker and Murphy, 1988). As in the standard model, rational addicts are forward-looking. Thus, one test of the rational addiction model is to see if individuals reduce current levels of consumption of addictive goods in response to anticipated price changes (Becker, Grossman, and Murphy, 1994). Critics have pointed out that this prediction of forward-looking behavior is not unique to the standard model. In fact, it accords with "all plausible psychological theories of people's errors," with the exception of a model of complete myopia (Rabin, 2013). Evidence in favor of the standard model may be weaker than its proponents have claimed.

Another line of criticism has been leveled specifically against the rational addiction model. Critics have pointed out that supportive evidence for this model has been drawn largely from analyses that look for positive consumption levels such that a person would choose to remain addicted, rather than scrutinizing the decision to become addicted in the first place (Rabin, 2013). If addictions are as powerful as suggested, then the decision to become addicted is of far greater relevance to understanding the addictive process. As above, the evidence base is not fully persuasive.

SELF-CONTROL PROBLEMS

A bias for the present

Self-control problems may be thought of as intra-personal conflicts that pit a present self against a stream of future selves. For example, Stu feels an urge to binge on the chocolate ice cream in his freezer, but he knows that his future self will regret it. Moreover, had we asked Stu yesterday about whether he wanted his present self to eat the ice cream, he would have said no. Stu's various incarnations are in conflict, a clear indication that his preferences are time-inconsistent. Another way to put it is that eating the ice cream would represent a preference reversal. Yesterday he would have preferred for his today's self to avoid the ice cream, but today the chocolate Sirens are calling.

Economists refer to self-control problems as *present-biased preferences*, or simply present bias, because they are the outcome of placing too much weight on present costs and benefits and too little on future costs and benefits. All self-control problems involve this feature of revising our consumption plans. We make plans for tomorrow that we revise once tomorrow comes. There are two ways that we tend to revise our plans. First, consider the case of "investment goods" that have an immediate cost and a future benefit. These are healthy goods that we wish to consume but lack the motivation or willpower to do so, such as physical exercise and preventive services such as cancer screening. A present-biased person will tend to delay the consumption of investment goods, going to the gym and getting cancer screening too infrequently. The gym members plan to work out but revise those plans when it comes time to peel themselves off the couch. For most people, working out is a chore. We would much rather loaf on the couch than deal with the immediate unpleasantness that the workout brings, and the potential benefits involving weight loss and improved health are far in the future. Our present self places far more importance on those immediate costs than the long-run benefits, and so we skip the gym workout today in hopes that we will make a different choice tomorrow. Often, we make the same choice tomorrow and delay the activity yet again.

The second scenario of revising our plans involves "leisure goods" that have an immediate benefit and a future cost. These are unhealthy goods that we wish to consume in moderation or not at all but then overindulge, such as junk food, recreational drugs, and risky sexual behavior. Present-biased individuals will over-consume leisure goods because of the urge to realize the immediate gratification. We want to smoke cigarettes today and cut back or quit tomorrow. When tomorrow comes, we revise our plan and consume as we did the day before. In this way, we may maintain a bad habit indefinitely, even though we would prefer to kick it.

There are a number of telltale signs that self-control problems are commonplace in the general population. Rational choice theory may be able to account for some of

these aspects of behavior, e.g., Orphanides and Zervos' (1995) attempt to incorporate regret. Taken as a whole, however, these patterns are highly suggestive that models that allow for self-control problems constitute a more parsimonious explanation.

First, individuals often express a desire not to partake in the unhealthy behaviors in which they then engage, and they often make unsuccessful attempts to quit or reduce consumption of unhealthy goods. For example, 69% of current U.S. smokers reported in 2010 that they would like to stop smoking completely (Malarcher et al., 2011). More than half of smokers made a quit attempt in the prior year, but only 6% of smokers quit successfully (Malarcher et al., 2011). Failed quit attempts can exact hefty costs in terms of cravings, withdrawal symptoms, and a bruised ego. Most individuals would not voluntarily subject themselves to this pain unless they really wanted to succeed. Moreover, individuals frequently make repeated failed attempts to change behavior. In one typical study, a sample of smokers had a mean of four past failures (Zhou et al., 2009). This pattern of behavior would be a highly inefficient path to maximizing utility.

Second, individuals who succeed in changing their behavior often succumb to high rates of recidivism. Studies find that as many as 40 to 60% of patients treated for drug dependence return to active use within a year of discharge from treatment (McClellan et al., 2000). Addictions in general are characterized by a chronic risk of relapse.

A third piece of evidence that self-control problems are common is the prevalence of regret that individuals express about past behavior. For example, 90% of smokers in Australia, Canada, the U.K., and the U.S. agree or strongly agree that they would not have started smoking if they could do it over again (Fong et al., 2004). It is no wonder that individuals often turn to professionals for help, as seen by the proliferation of behavioral and cognitive therapies to modify lifestyles.

Beyond the existence of anomalous behavioral patterns, a large body of research has directly elicited the time preferences of individuals and found them to be dynamically inconsistent. This research consistently shows the tendency of individuals to discount delayed rewards according to a hyperbolic function, in contrast to the exponential function assumed by standard economic theory (Ainslie, 1992; Frederick, Loewenstein, and O'Donoghue, 2002). Relative to the exponential function, the hyperbolic function has a steeper decline over small delays but levels out into a smaller decline over long delays. For example, in trading off a tempting snack now versus tomorrow (small delay) and tomorrow versus next week (long delay), a person would be expected to heavily discount the value of the snack between now and tomorrow but to modestly discount the value of the snack between tomorrow and next week. Figure 1 depicts the shape of these discounting functions, as well as the quasi-hyperbolic function discussed in the next section.

Figure 1. Discount functions

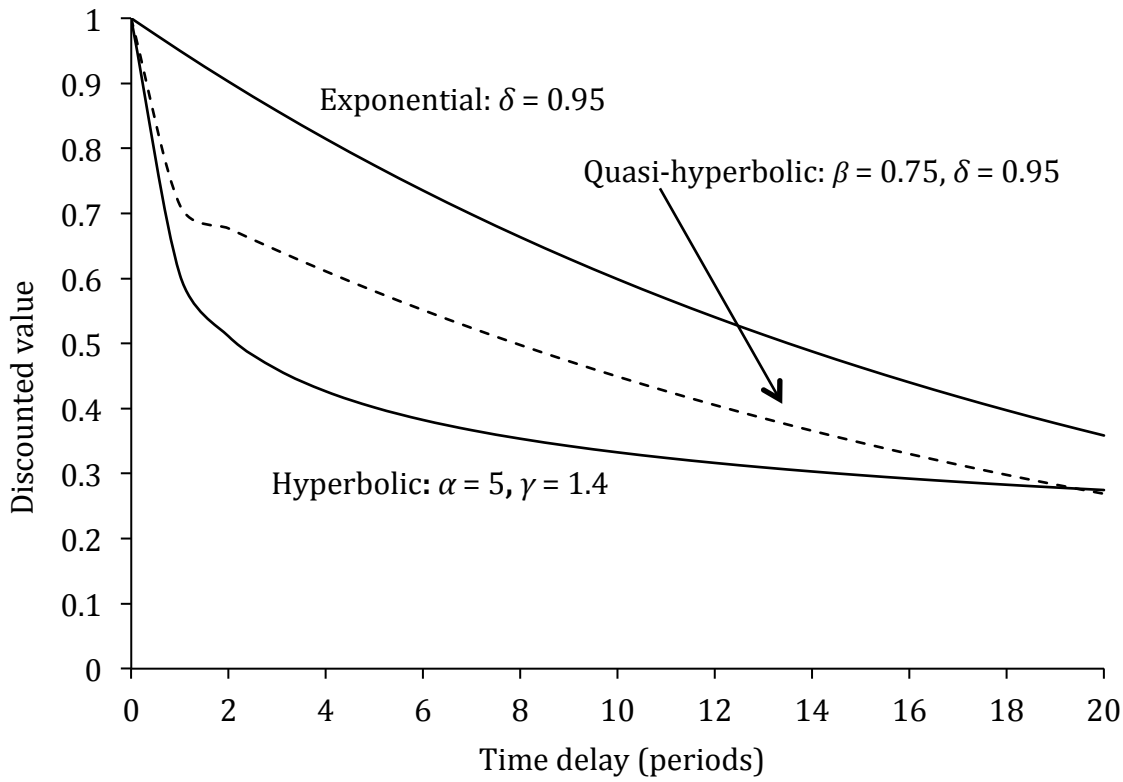

Note: The hyperbolic discount function is: $(1 + \alpha t)^{-\gamma/\alpha}$, as used in Loewenstein and Prelec (1992).

Dozens of empirical studies have linked hyperbolic discounting to health behaviors, such as tobacco use, alcohol use, and illicit drug use (Chabris, Laibson, and Shuldt,, 2008). These studies typically calculate discount rates by asking individuals about their willingness to accept varying amounts of money in exchange for varying delays of gratification. To take one example, Kirby, Petry, and Bickel (1999) observe the responses of heroin users and matched controls to offers of smaller immediate rewards ($11 to $80) and larger delayed rewards ($25 to $85). They show that a hyperbolic discounting function provides a good fit for the data, and heroin users hold discount rates about twice as large as the control group. Similar patterns have been found for users of nicotine, alcohol, cocaine, and other substances.

A model of self-control problems

In Section 2.2, we presented a utility function that assumes individuals are exponential discounters. Economists have suggested other discount functions that do not embed an assumption of time-consistency. R. H. Strotz (1956) pointed out that any discount function other than an exponential function could produce time-inconsistent preferences, although he did not advocate any specific functional form.

David Laibson (1997), drawing on work by Phelps and Pollak (1968), proposed an easy-to-use model of time-inconsistent preferences that has a quasi-hyperbolic discount function. Much like the hyperbolic discount function advocated by Ainslie (1992) and others, the Laibson model allowed for steeper discounting in the short term than in the long term:

$$U = u_0 + \beta\delta u_1 + \beta\delta^2 u_2 + \beta\delta^3 u_3 + \cdots + \beta\delta^T u_T$$
$$= u_0 + \beta \sum_{t=1}^{T} \delta^t u_t \qquad (4)$$

The only difference between Laibson's β-δ model in Equation (4) and the standard model in Equation (1) is that the β-δ model embeds an additional discount factor $\beta \leq 1$ on all future utility. In fact, the standard model is a special case of the β-δ model where $\beta = 1$. If $\beta = 1$, then Equation (4) simplifies to Equation (1), and the person becomes an exponential discounter with time-consistent preferences. If $\beta < 1$, then the person gives less weight to all future time periods relative to the present, and is a hyperbolic discounter with time-inconsistent preferences.

To see this, consider that a person in the β-δ model discounts utility tomorrow (period 1) relative to today (period 0) by a factor of $\beta\delta$. Now, consider that a person discounts any two consecutive future periods, such as tomorrow (period 1) versus the day after tomorrow (period 2), by a factor of: $\beta\delta^2/\beta\delta = \delta$. This is the same discount factor as in the standard model. In other words, the model predicts that a person is not subject to self-control problems when comparing the future to the more distant future. Only in the present does the person feel the pull of temptation. Thus, we want to eat candy today but avoid it tomorrow. However, when tomorrow comes, we again indulge ourselves and decide to avoid candy the day after tomorrow. Once the day after tomorrow comes, we again revise our plans. This is the essence of time-inconsistent preferences for a hyperbolic discounter.

As an aside, it is as yet unclear the extent to which self-control problems are domain-specific. Does a person's propensity to succumb to glazed donuts imply that the person will have poor self-control at work or at the pub? We do not know whether we should speak of individuals having a self-control parameter β or many βs, one for each type of action. We hope that future research will clarify this point.

Sophistication vs. naiveté

An important distinction is to be made between individuals who are aware of their self-control problems and individuals who are not. Behavioral economists refer to those who are self-aware as "sophisticated individuals" and those who are not as "naïve individuals" (O'Donoghue and Rabin, 1999a, 2001). In practice, many individuals do not fit neatly into either category. Self-awareness might be better conceived of as a continuum with full sophistication and full naïveté as the poles.

Most of us fall somewhere in the middle, partially aware of our self-control problems; we recognize that we lacked self-control today but nevertheless remain overly optimistic about our ability to show self-control in the future. A present-biased person's place on this continuum of self-awareness has major implications for how he deals with a self-control problem.

Sophisticated individuals may give in to temptation, but only for a limited period of time. Eventually, a sophisticate will search for strategies that precommit herself to following through on a goal. Consider a sophisticate who is trying to avoid over-consuming alcohol, a leisure good that has a present benefit and a future cost. She may pursue a few different strategies that precommit her to lower consumption of alcohol. The person may try to remove environmental cues that trigger drinking. For example, she may avoid hanging out with drinking buddies or avoid spending time in pubs. Or, she may try to diminish the pleasure she gets from alcohol. An extreme example would be the use of disulfiram, a drug that makes a person vomit when taken with alcohol. Needless to say, the unpleasant effects of the drug may be a powerful motivator for decreasing alcohol use.

In contrast to sophisticates, individuals who are naïve about their self-control problems will not recognize the need to precommit themselves to a path. In the extreme case, an aspiring gym-goer may delay working out day after day, year after year, always planning to go tomorrow. Likewise, the naïve dieter will perpetually plan to start a diet tomorrow, that is, until tomorrow comes and the grocery store has a sale on chocolate bars. If these individuals truly believe that they have the willpower needed to exercise or diet, they will fail to adopt self-control strategies or to seek out external support and they will inevitably fail to reform their ways.

We can represent sophistication and naïveté in the β-δ model presented in the last section. Recall that $\beta < 1$ is an indicator of present bias. Let $\hat{\hat{\beta}}$ be a person's prediction of his or her future self-control, that is, a prediction of β. A fully sophisticated person, who perfectly predicts the degree to which she will lack self-control, will have $\beta = \hat{\hat{\beta}} < 1$. A fully naïve person, who predicts that he will show total self-control, will have $\beta < \hat{\hat{\beta}} = 1$. A partially naïve person underestimates her future self-control, such that $\beta < \hat{\hat{\beta}} < 1$.

One conundrum of the sophistication-naïveté distinction is why individuals would not learn about their degree of present bias over time. After breaking a diet 100 times, wouldn't a dieter have learned by the 101st attempt? Moreover, can we educate naïfs about their self-control problems, for example by offering them feedback about past behavior, in order to increase their level of sophistication? More research is needed to understand the degree to which learning about self-control occurs over time. If we examine our own experiences, we find many instances in which our self-control problems have persisted over long stretches of time. Thus, we must theorize some reason that learning is incomplete, namely a reason that individuals do not adjust their beliefs to take self-control problems into

account. One hint may lie in what psychologists refer to as a self-serving bias. We tend to blame our failures on external factors outside of our control, on the situation in which we were placed rather than on our failings. We binged on snacks at the holiday party, but it was the party-planner's fault, not due to our own self-control. If we perceive each situation to be sufficiently different from the last situation—the holiday party is only once a year—then we may be able to convince ourselves that such instances are not learning opportunities.

Field evidence of self-control problems

The empirical literature on present bias has multiplied in recent years. In this section, we present several lines of economic research that highlight individuals' self-control problems related to health.

Exercise. Exercise is a prime example of a health application in which many individuals fail to meet their own goals. DellaVigna and Malmendier (2006) show that individuals make sub-optimal decisions about gym attendance. Gym members often purchase monthly passes yet attend the gym so infrequently that per-visit passes would be a cheaper alternative. Such inefficient behavior is incongruous with a standard model, although it can be reconciled with a model of present bias. The authors conclude that sophisticates who recognize their self-control issues may purchase monthly passes as a precommitment to exercise more often.

Acland and Levy (2014) show that infrequent gym visits may also involve a high degree of naïveté about present bias. As part of a field experiment, participants were eligible to win cash in exchange for meeting attendance targets in a randomly assigned week. Comparing participants' predictions about attendance and actual attendance, the authors find that participants are highly overoptimistic (i.e., naïve) about future gym attendance. Participants expect their future selves to be two-thirds less present-biased than they actually are now.

Tobacco use. Several economists have noted the link between present bias and the consumption of addictive goods. Satisfying one's immediate gratification today for an addictive good, such as tobacco, makes a person more likely to become addicted tomorrow than that person would have preferred in advance.

Gruber and Köszegi (2001) reassess the evidence in favor of the rational addiction model, such as the work of Becker, Grossman, and Murphy (1994) based on the forward-looking behavior of smokers. They show that observed smoking patterns fit a model in which smokers have present-biased preferences. In particular, present-biased smokers, like "rational" smokers, would respond now to expected future increases in excise taxes, although they may respond less than the rational addiction theory would predict.

Another piece of evidence comes from subjective ratings of happiness before and after the passage of tobacco control legislation. Gruber and Mullainathan (2005)

find that smokers report having greater life satisfaction after the regulations take effect, as would be the case if smokers view smoking as a bad habit that they regret. In this context, the tobacco control laws function as a precommitment that aids smokers in following through with their own plans to quit smoking.

Levy (2010) finds more direct evidence on the present bias of smokers. He is able to separate the short-run and long-run discount factors, β and δ, by using short-run and long-run changes to cigarette prices. Price variation due to taxes represents a long-run price change, and fluctuation in tobacco leaf prices represents a short-run change. Assuming that smokers are sophisticated about their present bias, Levy finds that on average smokers exhibit a moderate level of present bias that leads many to start a welfare-reducing, lifelong smoking habit.

Default effects. Another line of evidence comes from the potency of default options in influencing individuals' decision-making. One of the most notable examples of default effects is organ donation. Some countries require citizens to register to become organ donors (opt-in system), whereas others require citizens to register if they do not want to donate their organs (opt-out system). Johnson and Goldstein (2003) find that more than 85% of citizens donate in countries that have an opt-out system, and over 99% in most of opt-out countries, as compared to 4% to 28% under an opt-in system. At least one aspect of these sizable default effects is likely the role of present bias.[3] Opting into the system involves time and mental effort to complete and submit the requisite forms. This process may be unpleasant, as individuals are asked to contemplate their wishes if they die. As we have discussed, situations that carry upfront costs and long-term benefits are apt to be put off, sometimes indefinitely.

Default effects have also been found for vaccination (Chapman et al., 2010), advance directives (Halpern et al., 2013), and HIV testing (Montoy, Dow, Kaplan, 2014). Moreover, they likely apply to many other aspects of health, such as choice of a health insurance plan, contributions to health savings accounts, health worker flu shots, and treatment decisions like the removal of a catheter after 72 hours in order to avoid urinary tract infections. Halpern, Ubel, and Asch (2007) discuss many other examples of default policies that may improve health.

Payday effects. Many health-related programs dispense regular cash or in-kind benefits. Examples include food assistance programs and conditional cash transfer programs that make payments conditional on criteria such as regular check-ups or vaccinations. In many programs, the purchasing behavior of recipients shows signs of a regular cycle, in which recipients make larger or more frequent purchases around the time of receipt of income but then run out of money by the end of the

[3] Other factors contributing to default effects include a lack of an understanding of the choice situation, a perception that the default is an endorsement of a certain choice, the complexity of the choice leading to delay, and inattention to non-default options.

cycle. The consumption cycle is highly suggestive that individuals have a short-run impatience, or present bias.

A good example is the monthly food stamp cycle in which some families that receive monthly benefits through the Supplemental Nutrition Assistance Program (SNAP) consume relatively more food and calories at the start of the month and run out of food by the end of the month (Wilde and Ranney, 2000; Shapiro, 2005, Hastings and Washington, 2010). Food insecurity from this cycling may put diabetic SNAP recipients at higher risk of end-of-month hypoglycemia and hospital admissions (Seligman et al., 2014). This monthly cycling cannot be explained under the standard model, although it is quite consistent with behavior under a model of quasi-hyperbolic discounting.

Another consequence of payday effects is that individuals may be more likely to purchase "temptation goods," such as alcohol and sweets immediately after the receipt of income. Dasso and Fernandez (2013) find that recipients of a conditional cash transfer in Peru are more likely to spend the income on sweets, soft drinks, and meals in restaurants, relative to non-recipients and relative to themselves during other times of the month. This behavior violates the assumption from the standard model—known as the permanent income hypothesis—that receiving previously anticipated sources of income should not affect a household's consumption.

Deadline effects. Another line of evidence, largely unexplored for health-related decisions, is the way in which individuals respond to deadlines or time-limited offers. Individuals will often prefer to defer a costly action until the future, although individuals who are overoptimistic (i.e., somewhat naïve) about their future self-control may procrastinate. Setting a deadline or making a time-limited offer, if sufficiently large to overcome the immediate cost, may be able to prevent present-biased individuals from procrastinating on the action. Duflo, Kremer, and Robinson (2011) offer a time-limited discount on fertilizer to Kenyan farmers right after the harvest when they are flush with cash, and find that the deadline significantly helps to ward off procrastination. Ariely and Wertenbroch (2002) show that deadlines help workers not only to complete a task but also to do a more thorough job. Moreover, they find that evenly-spaced, externally imposed deadlines are more effective than self-imposed deadlines. While many health settings feature deadlines, such as the open enrollment period for selecting a health insurance plan, these aspects have not been studied in depth.

Precommitment. Perhaps the most notable area of research on self-control problems is individuals' demand for precommitment. These are strategies that make it more costly for a person to deviate tomorrow from a person's preferred choices today. Individuals need to be at least somewhat self-aware (at least partially sophisticated) about their future self-control in order to seek out precommitment strategies. For example, a person who recognizes her weakness for ice cream may avoid buying any at the grocery store, thereby making it more difficult to binge when she craves ice cream tomorrow. In a world of rational choice,

that is, in a world of perfect self-control, individuals would have no need for precommitment. We take up further discussion of precommitment in Section 6.2.

Microfoundations of self-control problems

As yet, we have not discussed in depth the biological and psychological underpinnings of self-control problems, except to say that they involve intertemporal tradeoffs. Behavioral economists lean on two overlapping frameworks for understanding the deeper underpinnings, or microfoundations, of self-control problems, as captured by hyperbolic and quasi-hyperbolic discounting. The first framework looks at the neurological foundations of self-control problems. This burgeoning field, known as neuroeconomics, points to an interplay between multiple systems in the brain that process information on rewards. The second framework focuses on the psychological foundations of self-control problems, notably the dual cognitive processes that lead to preference reversals.

Neuroeconomics

Multiple neurological systems process information on rewards, each of which handles time delays differently. The mesolimbic dopamine system is "impatient"; the fronto-parietal system is "patient." Although the neurobiological evidence is still subject to multiple interpretations (Sellito, Ciaramelli, and di Pellegrino, 2011), one view is that the balance between levels of activation of these two systems determines the degree of resulting patience or impulsivity. For example, McClure et al. (2004) use functional MRIs to measure brain activity during experiments in which subjects were offered monetary choices between smaller amounts of varying immediacy and larger delayed amounts. Regardless of the immediacy, there was activation of regions of the prefrontal cortex commonly associated with higher cognitive function such as computation. However, when choices involved more immediate rewards there was disproportionate activation of midbrain limbic structures associated with the dopamine system, i.e., areas that have been implicated in impulsive behavior. Research by Albrecht et al. (2011) building on this study further found immediate awards disproportionately associated with increased activity in the anterior cingulate cortex, which they interestingly note has also been found to be activated in risky decision problems that involve large gains. Relatedly, studies from patients with brain lesions that reduce emotional responsiveness find that those patients display lower levels of myopic loss aversion (Shiv et al., 2005). This is consistent with the hypothesis that individuals may make decisions that better correspond to standard rational choice models when the parts of the brain engaging in executive function have less competition from emotional centers.

Dual-Process Models

Psychologists have advanced a number of frameworks describing cognitive processes in conflict. While the shorthand used to describe the processes has varied

by author, the concepts are relatively stable. Humans have one cognitive process that is impulsive, reflexive, impatient, and irrational and another that is deliberative, analytical, patient, and rational. For example, Shiffrin and Schneider (1977) refer to automatic and controlled processing; Epstein (1994) to the experiential system and the rational system; Sloman to associative thought and rule-based reasoning; Metcalfe and Mischel (1999) to hot and cool cognitive systems; and Kahneman (2011) simply to System 1 and System 2. The conflict between these processes waylays a person from any premeditated plans and induces self-control problems.

In recent decades, behavioral economists have incorporated dual-process thinking into their own models under the broad category of "dual-self models." These assume that preferences can be represented by multiple selves within the person, whose interests are only partially aligned. This tradition dates back to Adam Smith (1759), who recognized in *The Theory of Moral Sentiments* the internal struggle between the "passions" and the "impartial spectator." Only since the emergence of behavioral economics have economists revived this work. Here, we introduce a handful of dual-self models to provide a flavor for their distinguishing and overlapping features and implications. This discussion draws heavily on a review of precommitment by Bryan, Karlan, and Nelson (2010).

Thaler and Shefrin (1981) propose a dual-self model comprised of a planner and a doer. The planner is a farsighted agent who wishes to maximize the person's lifetime utility but does not actually make any consumption decisions; rather, she derives utility from the choices of the doer. The doer lives for one period and is completely selfish and myopic. At times, the planner may try to modify the preferences or incentives of the doer. For example, a planner on a diet may increase monitoring of the doer by counting calories. Keeping track of the calories acts as a tax on the doer's decision to break the diet. At other times, the planner may wish to constrain the choices of the doer by imposing certain rules. These may be flexible rules, such as a dieter who aims to cut refined sugar from her diet, or stringent rules, such as checking herself into a weight loss clinic that will enforce a strict diet. We come back to these rules in our discussion of precommitment in Section 6.2.

There are several variants of the planner-doer model. One difference arises due to different assumptions about the short-run self. Fudenberg and Levine (2005) assume that the short-run self is myopic. Bernheim and Rangel (2004) describe a short-run self, an addict, who is helpless to cue-triggered impulses. Each dual-self model implies that individuals have different abilities to control impulses. The long-run self may be able to incur a cost in exchange for exerting control over the short-run self (Thaler and Shefrin, 1981; Fudenberg and Levine, 2006) or may face a random shock that determines whether the long-run self or short-run self has complete cognitive control (Bernheim and Rangel, 2004). In each case, the models predict that precommitment strategies would be helpful to the long-run self.

A debate in the literature involves whether dual-self models are consistent with an alternative model of self-control problems developed by Gul and Pesendorfer

(2001). In their temptation model, individuals have preferences regarding menus, or choice sets, of items. The canonical example is a customer at a restaurant who is choosing between, say, a hamburger and a salad. The person receives a certain level of utility from choosing from a menu where salad is the only entrée listed. However, once the hamburger is also on the menu, the salad provides less satisfaction, as the self-control needed to choose the salad and overcome the tempting hamburger is costly. More generally, a person's welfare may diminish if a tempting option is added to the choice set. Unlike the β-δ model, the temptation model implies that choices are always consistent with preferences. A person who gives in to temptation does so because the cost of temptation outweighs the cost of exerting self-control. In the β-δ model, utility depends only on actions; costly self-control does not factor into a person's utility function. Yet, as in the β-δ model, Gul and Pesendorfer's model predicts that a person will seek out precommitments to remove tempting options from the person's future choice set.

MISTAKEN BELIEFS

In this section, we discuss situations in which a person's beliefs deviate from rational expectations. We consider two forms of mistaken beliefs: 1) mispredictions about future preferences, and 2) overconfidence.

Mispredicting preferences

An old adage is never to shop on an empty stomach. In a classic experiment, Read and van Leeuwen (1998) studied the food choices of office workers in different hunger states, either when hungry late in the afternoon or when sated immediately after lunch. Subjects in each hunger state were asked to order a snack from a menu of healthy snacks (e.g., apples and bananas) and unhealthy snacks (e.g., candy bars and potato chips) one week in advance. Hungry subjects would be expected to have less self-control and to choose the unhealthy option. Some subjects were told that the food would be delivered late in the afternoon when they were hungry, while others were told that the food would be delivered right after lunch when they were sated. Thus, in making the choice, subjects had to grapple with their current hunger state and their predicted hunger state in one's week time. The study revealed that subjects correctly anticipated being less likely to want an unhealthy snack in a future state of satiety (last row of Table 1), although subjects allowed their current state of hunger to heavily influence this decision. In particular, subjects projected their current state of hunger onto their predicted future hunger. Whereas 78% of hungry subjects ordered an unhealthy snack to be delivered when hungry in a week, only 42% of sated subjects did so. The same holds for those predicting snack preference for a future state of satiety: hungry subjects were twice as likely as sated subjects to prefer the unhealthy snack to be delivered when sated in a week.

Table 1. Percentage of subjects who chose an unhealthy snack, by current and future hunger

	Hungry in a week	Sated in a week
Hungry now	78%	56%
Sated now	42%	26%
Total	60%	46%

Source: Read and van Leeuwen (1998)

The tendency to project one's current state onto one's predictions for the future is a general phenomenon called *projection bias* (Loewenstein, O'Donoghue, and Rabin, 2003). In the example above in which subjects are asked to predict their future hunger state, they are making a mistake by projecting their current state onto points in the future when their current state is no longer relevant. An important case of projection bias occurs when a person's state of being is affected by visceral factors, such as emotions (e.g., fear, anger), drive states (e.g., hunger, sexual desire), and feeling states (e.g., pain). Loewenstein (2005) refers to these visceral factors as affectively "hot" states, as opposed to affectively "cold" states when a person is not affectively aroused.

Loewenstein describes two general forms of projection bias involving visceral factors. When visceral factors grip our minds, we fail to appreciate the extent to which the affective state influences our preferences and behavior. This leads to a *hot-to-cold empathy gap*, in which we underestimate how our preferences will change once we exit the temporary hot state. Put differently, we overestimate how stable our preferences are when in a hot state. We saw this in the Read and van Leeuwen article; hungry workers had trouble divorcing their current hunger from their snack choice for a future state of satiety. We may also suffer from a *cold-to-hot empathy gap*. In our more sober moments, we do not predict the extent to which our decision-making will be distorted in an affectively hot state. For example, on a full stomach, we may recognize that hunger would push us toward choosing an unhealthy snack, but we underestimate the degree to which this force will drive us.

Many important health decisions occur once we have entered a hot state. The undue influence that visceral factors exert on our decision-making process often leads to errors. Below, we sample from the evidence on projection bias and health.

Addiction. Addiction is one of the most important arenas in which projection bias comes into play. Addicts in a state of craving may underestimate how good they would feel if unaddicted. Addicts may exaggerate how long the pain of quitting will last. Moreover, non-addicts may not fully appreciate the motivational force of addiction, contributing to initial decisions to experiment with drugs. All of these forces could contribute to the development of harmful addictions.

Several studies provide evidence on the projection bias of addicts. Giordano et al. (2002) and Badger et al. (2007) assessed the time preferences of long-time heroin addicts who were receiving a less potent substitute, buprenorphine, to reduce cravings for heroin and alleviate opioid withdrawal symptoms. The addicts typically received a single maintenance dose per day. The researchers offered the addicts a choice between different amounts of money and a second dose of buprenorphine, which is still pleasurable, to be delivered after the initial dose. Some would receive the dose that day and others in five days. The key experimental manipulation was that half of the subjects were asked when "deprived," two hours before receiving their initial daily dose, and half were asked when "sated," right after receiving their initial daily dose. The authors found that, for doses to be delivered that day, subjects were willing to pay an average of $75 when currently deprived and $50 when currently sated. For those asked about a second dose to be delivered in five days, the authors found a similar pattern; deprived subjects were willing to pay $60 for the second dose and sated subjects were willing to pay $35. Note that all subjects were to receive the second dose when in the same craving state, i.e., right after the initial dose. Thus, these findings reveal that even experienced addicts did not fully appreciate how much their preferences for drugs would change after the cravings subside.

Levy (2010) studied projection bias in cigarette smokers. He finds that younger smokers do not fully appreciate their likelihood of becoming addicted in the future, underestimating by 40% the degree to which their preferences for being addicted will change once addicted.

Good habits. It is also the case that we may not realize the extent to which we would enjoy developing good habits. For example, we may under-appreciate how much we would enjoy including more vegetables in our diet. Acland and Levy (2013) test whether a group of college students predict how much they would enjoy developing a habit for attending the gym. The authors paid non-gym attenders assigned to the treatment group to attend the gym at least twice a week for a month, thereby developing a "gym habit." The authors elicited predictions from subjects about their future gym attendance in order to compare the treatment group's valuation of going to the gym before and after developing the habit, relative to that of the control group. The authors found that subjects displayed near-total myopia—that is, projection bias—regarding the enjoyment they received from going to the gym after forming the habit. This suggests that researchers might place more emphasis on interventions that develop good habits, as they can help individuals to realize benefits that they otherwise would ignore.

Overconfidence

Individuals consistently overestimate their ability to perform tasks. A corollary to this statement is that individuals often see themselves as being less at risk of harm

compared with other people. Psychologists have termed this phenomenon *optimism bias*.

The origins of overconfidence include a combination of cognitive and motivational processes (Armor and Taylor, 1998). Cognitively, individuals rely on scenarios for how the future will unfold. These scenarios tend to be overly simplistic, to assume that individuals will engage in all actions necessary to obtain the outcome under consideration, and to neglect potential impediments such as situational factors and the actions of others. In comparing one's abilities to others, the person also may draw on reference groups that are not actually representative. Motivationally, individuals may maintain overoptimistic beliefs as part of a self-serving tendency in social comparisons. Affective states may also alter the degree of overconfidence a person holds.

Individuals display an unrealistic optimism about their susceptibility to a wide variety of health problems. Comparing individuals' objective health risk to their average perceived risk, individuals consistently report having a below-average perceived risk. This holds for the risk of lung cancer, tooth decay, ulcers, and a variety of other health risks (Weinstein, 1980, 1987). The result is that individuals may take more risks with their health than they realize and more than they would if better informed about the objective risk of a health behavior.

An important question is whether individuals accurately perceive their likelihood of becoming addicted to an addictive substance. The question bears on whether policymakers ought to educate the public or regulate addictive substances. The largest body of work in this area has focused on whether individuals accurately realize the addictive nature of tobacco use. While smokers seem to understand the health risks of smoking in general and in some cases may even overestimate those risks (Viscusi, 1992), that understanding does not necessarily translate into how smokers view their personal risk. Research shows that smokers underestimate their own personal risk relative to the risk of other smokers and nonsmokers (Slovic, 2001). In addition, smokers may perceive little or no risk from smoking individual cigarettes, even if they recognize the cumulative risks as high. Overconfident individuals, who underestimate their likelihood of becoming addicted, are likely to consume more cigarettes than they ever intended to consume.

In Section 3.3, we discussed overconfidence about an ability to show self-control. Individuals who are not fully aware of their self-control problems—that is, (partially) naïvely present-biased individuals—will overestimate the amount of self-control they will exert in the future. As discussed, naïveté can perpetuate a self-control problem over a long period of time and substantially harm the present-biased person.

Overconfidence would not be of great concern if it could be easily eliminated through education or some other means. However, the available research indicates that it is surprisingly difficult to "de-bias" a person who is overconfident (Weinstein

and Klein, 1995). More research is needed on the ways that researchers and policymakers might temper overconfident beliefs. One approach may lie in learning from situations in which individuals do not display overconfidence, such as for events that are perceived to be uncontrollable (Harris, 1996).

LIMITED ATTENTION AND MEMORY

Description

Individuals have a limited ability to attend to and process stimuli in their environment. Moreover, individuals lack the cognitive bandwidth to integrate all available information as part of some utility-maximizing procedure, as assumed by the standard model. Herbert Simon (1957) was one of the first to note these constraints, referring to them as the human capacity for *bounded rationality*. Subsequently, psychologists such as Kahneman and Tversky (1982) have catalogued some of the heuristics, or cognitive shortcuts, that individuals use to arrive at decisions when constrained by bounded rationality. We focus in this section on attention and memory as limited resources, as they are especially germane to intertemporal choices for health. Underweighting or wholly neglecting relevant information compromises our ability to make decisions in our own best interest. However, we acknowledge that other forms of nonstandard decision-making, such as a reliance on heuristics and a dependence on how choices are framed (e.g., as gains versus losses), are critical elements of decision theory more generally, not only for intertemporal choice.

When we make plans for the future, such as a weight loss plan, a critical prerequisite for success is that we remember the plan when the moment of truth arrives. In our busy lives, we might not remember the plan if we are buried under a pile of work or tending to a family crisis. Thus, memory poses one set of challenges. Even if we manage to remember the plan, we might not have an ability to attend to it amid multi-tasking or dealing with more pressing issues. It is important to recognize that attentional deficits are not always "irrational," especially if the acquisition of information needed to inform a decision is costly. We may want to lose weight, but it would require a heroic effort to make sense of the conflicting information in the popular media and scientific literature about the best way to lose weight. As such, we may be better off attending to other aspects of our diet, such as how often we should eat or in what quantity. The combination of memory and attentional failures are able to compromise even the best-laid plans.

A close cousin of attention is salience, the degree to which certain features of a choice or choice environment rise to the top of a person's mind. Some features appear more visible, or more salient, to a consumer while other features appear more opaque. Assume that we have settled on a plan to lose weight by limiting the amount of refined sugars in our diet. The nutritional content in our food is not

always obvious. Nutritional labels on packaged foods have not traditionally made it easy to discern nutritional content, and prepared foods make it even harder to monitor our dietary intake. The nutritional value of our food is often less salient than other features, such as the colors on the package or the brand name. (See Chapter 5 by Riis for further discussion of salient front-of-package nutritional labeling.)

In some cases, attributes of choices are shrouded from consumers. For example, the cost of visiting a clinic or a hospital in the U.S. is hidden from patients, meaning that patients do not typically take these costs into account when deciding where to seek care. Shrouding is often a deliberate profit-seeking strategy on the part of the companies that supply the good. In our example, the administrators of clinics and hospitals are able to charge higher prices to patients and health insurers if uninformed patients cannot direct their business toward more affordable facilities. In other cases, a person may overlook items when making plans. We have already noted that multi-tasking and bounded rationality hinder our ability to absorb relevant information. The end result is that limited attention and memory often prevent us from making health decisions that promote our well-being.

Applications

Researchers have studied limited attention and salience across a variety of health situations. In this section, we describe several areas of study.

Prices and taxes of health products. Chetty, Looney, and Kroft (2009) study whether consumers are attentive to the prices of beer. The authors use the fact that state sales taxes are not factored into the posted prices of goods but only assessed at the cash register, whereas excise taxes are included in posted prices. Thus, an excise tax may be considered more salient to the consumer than is a sales tax. The authors analyze cross-state differences in beer taxes and find that consumers are far more responsive to alcohol-specific excise taxes than to sales taxes. The authors rule out that customers are uninformed about the sales taxes. They run a survey of grocery shoppers indicating that they are informed about the sales tax. Rather, salience effects led consumers to underreact to taxes that are not included in posted prices. The broader lesson is that the salience of taxes and of incentives more generally may have a large impact on their effectiveness.

Reminders. Models of inattention make a key prediction about the responsiveness of individuals to reminders. Reminders can be used to draw attention to the future consequences of choices. Due in part to their simplicity, reminders have been an active area of investigation. The advancement of digital technology has made reminder systems increasingly easy to implement. It is now simple and inexpensive to send reminders to groups of people, and many individuals have access to phone and computer applications that can set reminders.

A natural application of reminders is to improve medication adherence. Vervloet et al. (2012) review the evidence on electronic reminders that are automatically sent to patients. Delivery mechanisms have included text message reminders and electronic reminder devices that alert users with lights or sounds. Most studies document short-run improvements in adherence, although the long-term impact of electronic reminders is not well established.

Several studies find that text message reminders are effective for promoting smoking abstinence (Free et al., 2011; Rodgers et al., 2005). This approach is easily transferrable to other attempts at behavior modification, including dieting and physical activity. Moreover, reminders could be effective for regularly scheduled activities, such as primary care visits, Pap smears, and colonoscopies.

Checklists. A growing movement of medical professionals has advocated for the use of checklists as a quality improvement intervention in medical settings (Haynes et al., 2009). Medical diagnoses and procedures routinely involve complicated and extensive lists of activities. In settings where clinicians are overwhelmed and systems of coordination are often inadequate, it is easy for mistakes to creep in. A simple checklist can be the difference between clinicians adhering to a protocol and committing gross negligence.

Planning prompts. One way to encourage individuals to follow through on a plan is to prompt them for details about how they intend to implement the plan. By prodding people to think through specific aspects of a task, it can reduce forgetfulness and prevent procrastination. Creating a plan forms an association between the future moment and the plan to be implemented. When the moment arrives, it cues the person to follow through on the plan.

Milkman et al. (2011, 2013) have found that simple planning prompts modestly increase preventive screening rates. Sending a planning prompt to get a flu shot by mail increased the share who received a flu shot by 4 percentage points above the 33% of control group members who got the shot, i.e., a 13% increase in relative terms. Similarly, a mailing that reminded people to get a colonoscopy increased the percentage of those who got the test by 1 percentage point above the 6% in the control group who got the test, i.e., a 16% relative increase. Like text message reminders, these are low-cost interventions that could easily be expanded population-wide.

IMPLICATIONS FOR POLICY AND BEHAVIOR

The role of policy

Under the standard model, there are two economic rationales for policy intervention into health-related behavior. First, policymakers may wish to correct a market imperfection that impedes the ability of individuals to fulfill their own preferences. Second, policymakers may wish to correct for costs that individuals impose on others, so-called negative externalities. Behavioral economic models suggest that there is also an important role for policymakers to address "internalities" that individuals impose on themselves from making sub-optimal choices. Several researchers have promoted policymaking that nudges individuals toward better choices without mandating change. This idea has been referred to as asymmetric paternalism or libertarian paternalism (Camerer et al., 2003; Thaler and Sunstein, 2003).

Many public health efforts involve the provision of information and education to individuals. Where imperfect information is a major contributor to poor decisions, these interventions may be effective at changing behavior. However, in instances where cognitive biases are responsible for unhealthy behavior, educational campaigns are not likely to achieve large gains. More robust behavioral interventions would be needed.

More generally, policymakers need to recognize the limits of human rationality. To take one example, U.S. lawmakers have taken a hard stance on drug use, including banning needle exchange programs in many states on the grounds that they send the wrong message to the public about the acceptability of drug use. A heroin addict who is due for a hit but lacks a clean needle is unlikely to think deeply about the consequences of injecting himself with a dirty needle. In the throes of a craving, he will feel an overpowering bias for the present and an inability to think beyond his preferences at that moment, not to mention potentially being overconfident about the likelihood of infection and being inattentive to where to access public health resources. Under these circumstances, the addict will often choose the dirty needle if a convenient alternative is not at hand. For any number of similar policies, there is a strong behavioral case for government action.

In the following subsections, we delve into some specific implications of behavioral economics for policy and behavior.

Commitment devices

A major contribution of behavioral economics to public health research has been the clearer direction it has provided about how to ameliorate self-control problems. Behavioral economists have highlighted the potential of self-management techniques, known as precommitments or commitment devices, in which individuals make it more costly for themselves to choose tempting options in the future.

Commitment devices have made their way outside of the ivory tower of academia. A number of websites, such as StickK.com and healthywage.com, have sprung up

offering willing customers the opportunity to bet on whether or not they can follow through on a health goal of the person's choice. The customers put their own money at risk, to be forfeited if they fail. To further raise the stakes, StickK.com even offers customers the option of donating forfeited money to an "anti-charity" of the person's choice, such as the National Rifle Association for a customer who believes in gun control. Some governments have also taken notice of these approaches. The U.K. government established an agency called the Behavioural Insights Team to investigate behavioral economic applications, and one of its efforts has involved testing the scale-up of commitment devices in a drug store chain in the U.K.

One way to classify commitment devices is as "soft commitments" or "hard commitments." The former are non-binding agreements to follow through on a goal, whereas the latter are binding agreements that typically carry a pre-specified penalty. We probe the difference between these two categories below.

Soft commitments. Soft commitments are promises that a person makes to oneself, or occasionally to others. As such, they tend to be self-enforcing contracts that a present self makes with his or her future self. An example might be a New Year's resolution where a person promises today to take some action tomorrow. Such personal rules may be helpful by boosting motivation and making certain types of behavior particularly salient. The challenge is that soft commitments can be difficult to enforce. In addition, the penalties, such as disappointment with oneself, may not be concrete enough to dissuade someone from breaking a personal rule.

Soft commitments may take different forms. They may involve attempts to avoid environmental cues that trigger temptation. Recovering alcoholics may avoid social events where alcohol is served, or dieters may avoid walking down the candy aisle of the grocery store. Soft commitments may also be personal rules that guide a set of behaviors, such as resolutions to smoke only after meals or to work out twice a week. Bénabou and Tirole (2004) argue that we create personal rules because we fear creating behavioral precedents that our future selves will follow. In the famous "marshmallow experiment," researchers showed that 4-year-olds who could delay immediate gratification in exchange for a larger reward—forgoing one marshmallow now in exchange for two later—experienced improved educational and social outcomes later in life (Mischel, Shoda, and Rodriguez, 1989). It is possible that these children had set a behavioral precedent for self-control that they followed throughout their lives.

To deter ourselves from allowing too many exceptions to a personal rule, we sometimes define "bright lines." Recovering alcoholics are often advised to avoid even a sip of alcohol. Dieters may avoid even a morsel of a forbidden food. Defining the boundaries of a rule may be useful for self-enforcement. A danger is that bright-line rules will cause a person to forgo large benefits, such as a dieter who gives up a turkey feast on Thanksgiving. More generally, personal rules sometimes lead us to over-regulate our behavior, occasionally turning a rule into an obsession or compulsion. Some people are able to carve out exceptions where it is acceptable to

deviate from the rule, putting a diet on hold for Thanksgiving, without self-signaling that a lack of willpower is acceptable.

Behavioral economists have focused most of their attention on hard commitments, although there has been some work on soft commitments. Khwaja, Silverman, and Sloan (2007) surveyed U.S. smokers and found that 81% had used a soft commitment device to limit tobacco consumption. Strategies included keeping busy by doing yardwork, banning smoking at home or at work, and avoiding other smokers. Some smokers also report "self-rationing," purchasing cigarettes in smaller quantities, such as by the pack rather than by the carton, in order to avoid consuming more than they prefer (Wertenbroch, 1998). Self-rationing may also occur for food choices, as in the case of consumers who only buy junk food in small quantities, knowing that it will limit the amount consumed during binges. Smokers may view tobacco control laws as a form of commitment device to limit tobacco consumption. Gruber and Mullainathan (2005) find that the passage of higher cigarette excise taxes increased the overall happiness of smokers, as would be predicted from a model of self-control problems but not from the standard model.

Hard commitments. These are binding commitments in which a pre-defined penalty for failing to reach a goal is enforced. Often called commitment contracts, these agreements typically put a person's own money at stake, which is contingent on meeting a health goal. Some permutations, called deposit contracts, ask people to deposit money on a regular basis as part of the monetary commitment. The money is typically forfeited if the person cannot reach the health goal. Other permutations earmark money for certain expense items, such as health emergencies, and the money cannot be withdrawn for other purposes. Though less tested, some contracts put a person's reputation at stake, as in the case of a public declaration to accomplish a task.

Hard commitments have been tested for a variety of health behaviors, including weight loss, smoking cessation, exercise, health savings, and use of anti-malarial bednets. Overall, the contracts have been effective at promoting behavior change across these domains. Weight loss is the one area where the results have not been consistently positive. The literature on commitment contracts is reviewed in greater detail in Chapter 8 of this volume by Lewis & Block. We briefly discuss several important features of these contracts here.

Commitment contracts have been promoted and criticized on a number of grounds. A major argument in their favor is their potential cost-effectiveness. Relative to clinical approaches, commitment contracts do not have the same reliance on skilled personnel and often use less expensive inputs. Commitment contracts may even generate net revenue, thanks to the forfeited deposits from users who fail to achieve their goal.

An outstanding question in the literature is the magnitude of demand for commitment contracts. Take-up rates of commitment products have varied greatly

across setting and application, though they have tended to be modest. Some survey data suggests that most individuals have unfavorable views toward the use of commitment contracts for health behavior change (Promberger et al., 2011). More research is needed to understand whether the approach is scalable and acceptable to the general public.

Commitment contracts are designed to be a "libertarian paternalistic" solution, leaving individuals without present bias unharmed, providing needed help to sophisticated present-biased agents, and leaving fully naïvely present-biased individuals unharmed but also unhelped. In practice, partial naïfs may fail to put enough money at stake to motivate themselves—under-committing—because they underestimate their self-control problem. As a result, these individuals may end up worse off than without the opportunity to commit, as at least they would have kept their money. For example, Giné, Karlan, and Zinman (2010) find that 66% of smokers who took up a deposit contract for smoking cessation failed to quit smoking. White, Dow, and Ruangrunghiranya (2013) supplemented a deposit contract with incentives and more than half of smokers still failed to quit. Researchers have yet to find a way around the issue of under-commitment.

A final issue is the inflexibility of hard commitments, which are binding by design. This comes at a cost to users. For example, if a person has a commitment savings account for health emergencies, what happens if a non-health emergency arises? Should the person be allowed to withdraw money from the account? It is as yet unclear if complete inflexibility is necessary for the success of health commitment contracts or whether a softer commitment might be just as effective.

Structure of incentives and programs

Next, we move to general principles that can diminish the negative impacts of cognitive biases. We believe that researchers and policymakers should attempt to incorporate these principles into programs and policies.

Self-control failures are caused by the temptation to capture immediate rewards and avoid immediate costs. One of the best ways for external incentives to compete with a tempting option is to make them available when the tempting option presents itself. A reward for weight loss will be more effective if offered at a weekly weigh-in than if offered at the end of a yearlong program. Moreover, the tendency to sharply discount any and all future consequences means that smaller immediate rewards can be at least as effective, if not more so, than larger rewards delivered in the future. Thornton (2005) shows that even the smallest rewards can substantially increase the willingness of individuals to learn their HIV status.

Public health practitioners often seek to change behavior over the long term. Rather than providing one big payoff at the end, a well-designed program should provide incremental incentives along the way. Doing so can increase the salience of the incentives and help stave off procrastination. O'Donoghue and Rabin (1999b) argue

that optimal incentives for procrastinators typically involve increasing the incentives as time passes. This is precisely the approach taken by behavioral researchers, who offered subjects an escalating reward schedule that paid drug users for abstinence in increasing amounts over time (Higgins et al., 1991, 2000). The approach has been successfully applied to several types of drugs.

The timing of incentives is only one factor of note. Another is the salience of incentives. In order for incentives to motivate, the potential recipients first need to recognize that the incentives are available. The more salient they are, the more likely they will change behavior. Some employers offer wellness programs that include rewards for completing certain activities, like a health risk assessment, or reaching certain goals, like blood pressure below some cutoff. These companies often provide the rewards as a discount on the person's health insurance premium, which only appears on the worker's paystub. A much more salient approach would be to hand cash directly to workers at the time it is earned.

Public health programs of all sorts can maximize their impact by harnessing the power of default options. In Section 3.4, we discussed the susceptibility of individuals to default options because the time and effort involved in making a choice can lead to procrastination. Default effects are accentuated by failures of attention and memory. We may not remember that we need to update our annual choice of health insurance plan or refill our psychiatric medications. This leads to a bias toward the status quo. Setting smart defaults can dramatically improve the behavior of distracted or present-biased individuals.

Optimal defaults should be set to encourage good health, as in the case of home delivery of chronic medications (Beshears, Choi, Laibson, and Madrian, 2014). Optimal defaults may involve requiring people to make an active choice, rather than simply falling back on the default option. Active choice ensures that the person's underlying preferences are respected and also has the virtue of ensuring that the default options are socially acceptable. Otherwise, the public health community faces the risk of backlash from pushing its priorities on an unwilling general public. In cases where one option is socially beneficial, such as organ donation, active-choice systems may be enhanced by framing the decision to highlight the benefits of the socially preferred option. For example, instead of asking "would you like to donate your organs?," we might ask "would you like to donate your organs in order to save lives?" (Keller et al., 2011) Default systems should set deadlines to ensure that individuals select a choice in a timely manner. Optimal defaults should typically avoid complex enrollment rules and procedures, restrictions on choices, and barriers to switching. The designer should aim to make it as easy as possible for individuals to implement their preferred choices.

Sin taxes

Taxation has been the preferred policy instrument of economists for regulating the consumption of risky health products, such as tobacco and alcohol. The empirical

literature has borne out the effectiveness of this approach. Countless studies have shown that individuals respond to the prices of risky health goods (e.g., IARC, 2011). Economists have traditionally argued that sin taxes are appropriate for recouping the external costs that consumers impose on society, i.e., the negative externalities. To the extent that consumers impose costs on themselves or their family, economists have traditionally argued that consumers already take those costs into account when choosing to consume the risky product, even if the product is addictive or harmful. According to this view, the government has no role in regulating products that do not produce externalities on net, aside from educating consumers about health risks and restricting the access of minors.[4]

Behavioral economists have enriched the traditional view of sin taxation. They have argued that sin taxes should also incorporate the costs from internalities, such as present bias. Cognitive biases prevent afflicted individuals from following through on their plans, and taxes may be able to help individuals to implement their long-run preferences.

Gruber and Köszegi (2001) apply the β-δ model to estimate the optimal level of tobacco taxes, including the cost of sophisticated present bias to smokers. Assuming a modest degree of present bias of $\beta = 0.9$, the authors calculate an optimal cigarette tax of $1 to $3 per pack to account for internalities from present bias. This estimate only includes the cost of premature death to the smoker. It does not include the costs of second-hand smoke, the long-run costs of low birth weight from mothers who smoke, smoking-related fires, the disutility to the smoker from smoking-related illness, and many other costs. Moreover, this estimate does not apply to naïvely present-biased smokers, for whom the optimal tax would likely be much higher.

Levy (2010) builds on the work of Gruber and Köszegi to calculate the optimal tax on smokers using an empirical estimate of present bias and adding the effect of projection bias to smokers. He estimates that the optimal tax rate would be $8 to $11 per pack, an order of magnitude larger than the current U.S. federal tax rate of roughly $1 per pack. Behavioral economists have only begun to apply these insights to other health goods.

A common concern of sin taxes is that they have the potential to be regressive, disproportionately harming low-income individuals. The potential regressivity of sin taxes becomes more complicated once behavioral factors are considered. Gruber and Köszegi (2004) find that tobacco taxes are substantially less regressive after factoring in that lower-income smokers are more sensitive to cigarette prices than are higher-income smokers. Cigarette taxes become less regressive as the degree of present bias increases in the population. That is because the self-control

[4] By "on net," we mean that the external costs outweigh the external benefits of the good. Indeed, at current levels of excise taxes in the US, analyses suggest that smokers are already subsidizing the health care costs of non-smokers, if restricting the analysis to only account for externalities.

benefits of tobacco taxation are largest for the lowest income groups such that on net they are better off after the taxation.

Differential impact on vulnerable populations

While cognitive biases affect all of us, there are good reasons to believe that certain vulnerable populations may be particularly affected. We focus on two vulnerable groups: low-income individuals, and adolescents and young adults. Identifying the groups at greatest risk is important for targeting interventions and, in the case of low-income individuals, for reducing health and economic disparities.

Low-income individuals. Low-income individuals may be hit hardest by cognitive biases. Avoiding mental mistakes requires time, effort, knowledge, and, at times, financial resources, all of which are on average more scarce among the poor. The poor face a staggering number of cognitive demands: finding and maintaining work, securing food, and navigating a complex, and frequently hostile, economic environment. With all of these competing demands, low-income individuals often have the least amount of cognitive bandwidth left over to allocate to intertemporal health choices. By virtue of their poverty, these individuals are also least likely to have access to decision aids that might keep cognitive biases in check.

There has been growing interest in the social determinants of cognitive biases. Several studies have established that cognitive performance decreases when a person is mentally taxed (Spears, 2011; Mani et al., 2013; Mullainathan and Shafir, 2013). The issue is not that the poor make bad choices because they are innately less capable, but rather that anyone who finds himself in a state of cognitive scarcity is prone to making mistakes. Poor people just tend to face severe mental strain on a more routine basis than the rest of us. A challenge for the field of behavioral economics is to develop interventions tailored to low-resource individuals and communities.

Adolescents and young adults. Researchers should also take special note of the effect of cognitive biases on adolescents and young adults. Neuroscientists have found that the development of the brain's limbic system continues throughout young adulthood, leading adolescents and young adults to have reduced cognitive control and increased sensitivity to rewards. This may account in part for the greater adoption of risky health behaviors in adolescence and young adulthood.

Many habits form early in life when individuals are least able to counter the pull of cognitive biases. For example, drug addicts typically first start to use in adolescence or young adulthood. These stages of life present critical time periods in which behavioral scientists might be able to intervene to influence the long-term trajectories of people's health. For example, Matthew Rabin (2013) has proposed that researchers experiment with age-based taxation, in which young people face higher tax rates on cigarettes or alcohol than the old.

CONCLUSION

Economists have traditionally relied on observed choices and behaviors to infer individuals' innate preferences for health and other aspects of their lives. Economists have gone so far as to refer to choices as "revealed preferences." In this chapter, we characterized ways that our intertemporal choices and our preferred plans can diverge. Cognitive errors undermine our plans and lead us to make sub-optimal choices about our health.

We focused on a handful of cognitive challenges: self-control problems, misprediction of future preferences, overconfidence, and limited attention and memory. Behavioral researchers have constructed a large body of work exploring each of these concepts. As the field of behavioral economics has grown, interest has spread to new topics of intertemporal choice, and the list of behavioral economic concepts will continue to grow in the coming years. Researchers have only begun to realize the potential of these existing and emerging insights for promoting health.

REFERENCES

Acland, D., & Levy, M. (2013). Naiveté, projection bias, and habit formation in gym attendance. Unpublished working paper. March 2013.

Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person*. New York, NY: Cambridge University Press.

Albrecht, K., Volz, G., Sutter, M., Laibson, D., & von Cramon, D. (2011). What is for me is not for you: Brain correlates of intertemporal choice for self and other. *Social Cognitive and Affective Neuroscience*, 6(2), 218-225.

Ariely, D., & Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13, 219-224.

Armor, D. A., & Taylor, S. E. (1998). Situated optimism: Specific outcome expectancies and self-regulation. *Advances in Experimental Social Psychology*, 30, 309–379.

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy*, 96(4), 675–700.

Badger, G. J., Bickel, W. K., Giordano, L. A., Jacobs, E. A., Loewenstein, G., & Marsch, L. (2007). Altered states: The impact of immediate craving on the valuation of current and future opioids. *Journal of Health Economics*, 26(5), 865-876.

Becker, G. S., Grossman, M., & Murphy, K. M. (1994). An empirical analysis of cigarette addiction. *American Economic Review*, 84(3): 396–418.

Bénabou, R., & Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, 112(4), 848–886.

Bernheim, B. D., & Rangel, A. (2004). Addiction and cue-triggered decision processes. *American Economic Review*, 94(5), 1558–1590.

Beshears, J., Choi, J. J., Laibson, D. I., & Madrian, B. C. (2014). Active choice and health care costs: Evidence from prescription drug home delivery. Unpublished working paper.

Bryan, G., Karlan, D., & Nelson, S. (2010). Commitment devices. *Annual Review of Economics*, 2(1), 671–698.

Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue, T., & Rabin, M. (2003). Regulation for Conservatives: Behavioral economics and the case for "Asymmetric Paternalism." *University of Pennsylvania Law Review*, 151(3), 1211–1254.

Chabris, C. F., Laibson, D. I., & Schuldt, J. P. (2008). Intertemporal choice. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*, 2nd Edition. New York, NY: Palgrave Macmillan.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4), 1145–1177.

Dasso, R., & Fernandez, F. (2013). Temptation goods and conditional cash transfers in Peru. Unpublished working paper. December 2013.

DellaVigna, S. (2009). Psychology and economics. *Journal of Economic Literature*, 47(2), 315–372.

DellaVigna, S., & Malmendier, U. (2006). Paying not to go to the gym. *American Economic Review*, 96(3), 694–719.

Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. *American Economic Review*, 101(6): 2350–90.

Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American Psychologist*, 49(8), 709–724.

Frederick, S, Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351–401.

Free, C., Knight, R., Robertson, S., Whittaker, R., Edwards, P., Zhou, W., Rodgers, A., Cairns, J., Kenward, M. G., & Roberts, I. (2011). Smoking cessation support delivered via mobile phone text messaging (txt2stop): A single-blind, randomised trial. *The Lancet*, 378, 49–55.

Fong, G. T., Hammond, D., Laux, F. L., Zanna, M. P., Cummings, K. M., Borland, R., & Ross, H. (2004). The near-universal experience of regret among smokers in four countries: Findings from the International Tobacco Control Policy Evaluation Survey. *Nicotine and Tobacco Research*, 6(Supp 3), S341–S351.

Giné, X., Karlan, D., & Zinman, J. (2010). Put your money where your butt is: A commitment contract for smoking cessation. *American Economic Journal: Applied Economics*, 2(4), 213–235.

Giordano, L. A., Bickel, W. K., Loewenstein, G., Jacobs, E. A., Marsch, L., & Badger, G. J. (2002). Mild opioid deprivation increases the degree that opioid-dependent outpatients discount delayed heroin and money. *Psychopharmacology*, 163, 174–82.

Gruber, J., & Köszegi, B. (2001). Is addiction rational? Theory and evidence. *Quarterly Journal of Economics*, 116(4), 1261–1303.

Gruber, J., & Köszegi, B. (2004). Tax incidence when individuals are time-inconsistent: The case of cigarette excise taxes. *Journal of Public Economics*, 88, 1959–1987.

Gruber, J., & Mullainathan, S. (2005). Do cigarette taxes make smokers happier? *The B.E. Journal of Economic Analysis & Policy*, 5(1).

Gul F., & Pesendorfer W. (2001). Temptation and self-control. *Econometrica*, 69(6):1403–35.

Halpern, S. D., Ubel, P. A., & Asch, D. A. (2007). Harnessing the power of default options to improve health care. *New England Journal of Medicine*, 357(13), 1340–1344.

Harris, P. (1996). Sufficient grounds for optimism? The relationship between perceived controllability and optimistic bias. *Journal of Social and Clinical Psychology*, 15, 9–52.

Hastings, J., & Washington, E. (2010). The first of the month effect: Consumer behavior and store responses. *American Economic Journal: Economic Policy*, 2(2), 142–162.

Haynes, A. B., Weiser, T. G., Berry, W. R., et al. (2009). A surgical safety checklist to reduce morbidity and mortality in a global population. *The New England Journal of Medicine*, 360, 491–499.

Higgins, S. T., Delaney, D. D., Budney, A. J., Bickel, W. K., Hughes, J. R., Foerg, F., & Fenwick, J. W. (1991). A behavioral approach to achieving initial cocaine abstinence. *American Journal of Psychiatry*, 148(9), 1218–1224.

Higgins, S. T., Wong, C. J., Badger, G. J., Haug Ogden, D. E., & Dantona, R. L. (2000). Contingent reinforcement increases cocaine abstinence during outpatient treatment and 1 year of follow-up. *Journal of Consulting and Clinical Psychology*, 68(1): 64-72.

International Agency for Research on Cancer (IARC) (2011). Effectiveness of tax and price policies for tobacco control. Chapter 5; 137–200.

Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, 302(5649), 1338–1339.

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus, and Giroux.

Kahneman, D., & Tversky, A. (2000). *Choices, values, and frames*. New York, NY: Cambridge University Press.

Kahneman, D., Tversky, A., & Slovic, P. (Eds.) (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.

Karlan, D., McConnell, M., Mullainathan, S., & Zinman, J. (2010). Getting to the top of mind: How reminders increase savings. National Bureau of Economic Research, Working Paper No. 16205.

Keller, P. A., Harlam, B., Loewenstein, G., & Volpp, K. G. (2011). Enhanced active choice: A new method to motivate behavior change." *Journal of Consumer Psychology*, 21(4), 376–383.

Khwaja, A., Silverman, D., & Sloan, F. (2007). Time preference, time discounting, and smoking decisions. *Journal of Health Economics*, 26, 927–949.

Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999) Heroin addicts discount delayed rewards at higher rates than non-drug using controls. *Journal of Experimental Psychology: General*, 128, 78–87.

Klabunde, C. N., Brown, M., Ballard-Barbash, R., White, M. C., Thompson, T., Plescia, M., & King, S. C. (2012). Cancer screening—United States, 2010. *MMWR*, 61(3), 41–45.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112, 443–77.

Levy, M. R. (2010). An empirical analysis of biases in cigarette addiction. Unpublished working paper. September 2010.

Loewenstein, G. (2005). Hot-cold empathy gaps and medical decision making, *Health Psychology* 24(4(Supp)), S49–S56.

Loewenstein, G., & O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics,* 118(4), 1209–1248.

Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: evidence and an interpretation. *Quarterly Journal of Economics*, 107, 573–97.

Malarcher, A., Dube, S., Shaw, L., Babb, S., & Kaufmann, R. (2011). Quitting smoking among adults — United States, 2001–2010. *MMWR*, 60(44), 1513–1519.

Mani, A., Mullainathan, S., Shafir, E., & Zhao, J. (2013). Poverty impedes cognitive function. *Science*, 341, 976–980.

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695), 503–507.

McLellan, A., Lewis, D. C., O'Brien, C. P., Kleber, H. D. (2000). Drug dependence, a chronic medical illness: Implications for treatment, insurance, and outcomes evaluation. *JAMA*, 184(13), 1689–95.

Metcalfe, J., & Mischel, W.. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, 106(1), 3–19.

Milkman, K. L., Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C., (2011). Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences*, 108, 10415–10420.

Milkman, K. L., Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C., (2013). Planning prompts as a means of increasing preventive screening rates. *Preventive Medicine*, 56, 92–93.

Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, 244, 933–938.

Montoy, J. C., Dow, W. H., & Kaplan, B. (2014). Active-choice and opt-out HIV screening: A randomized controlled trial. Unpublished working paper.

Mullainathan, S., & Shafir, E. (2013). *Scarcity: Why having too little means so much*. New York, NY: Times Books.

O'Donoghue, T., & Rabin, M. (1999a). Doing it now or later. *American Economic Review,* 89(1), 103-124.

O'Donoghue, T., & Rabin, M. (1999b). Incentives for procrastinators. *Quarterly Journal of Economics*, 114(3), 769–816.

O'Donoghue, T., & Rabin, M. (2001). Choice and procrastination. *Quarterly Journal of Economics,* 116, 121–60.

Orphanides, A., & Zervos, D. (1995). Rational addiction with learning and regret. *Journal of Political Economy*, 103(4), 739–58.

Osterberg L, Blaschke T. (2005). Adherence to medication. *New England Journal of Medicine*, 4; 353(5), 487-97.

Promberger, M., Brown, R. C. H., Ashcroft, R. E., & Marteau, T. M. (2011). Acceptability of financial incentives to improve health outcomes in UK and US samples. *Journal of Medical Ethics*, 37(11), 682–687.

Rabin, M. (2013). Healthy habits: Some thoughts on the role of public policy in healthful eating and exercise under limited rationality. In A. Oliver (Ed.), *Behavioural Public Policy* (pp. 115-139). New York, NY: Cambridge University Press.

Read, D., & van Leeuwen, B. (1998). Predicting Hunger: The effects of appetite and delay on choice. *Organizational Behavior and Human Decision Processes*, 76(2), 189–205.

Rodgers A., Corbett T., Bramley D., Riddell, T., Wills, M., Lin R., & Jones, M. (2005). Do u smoke after txt? Results of a randomised trial of smoking cessation using mobile phone text messaging. *Tobacco Control*, 14, 255–61.

Samuelson, P. (1937). A note on measurement of utility. *Review of Economic Studies*, 4, 155–61.

Schelling, T. C. (1980). The intimate contest for self-command. *Public Interest*, 60, 94-118.

Seligman, H. K., Bolger, A. F., Guzman, D., López, A., & Bibbins-Domingo, K. (2014). Exhaustion Of food budgets at month's end and hospital admissions for hypoglycemia. *Health Affairs*, 33(1), 116–123.

Sellito, M., Ciaramelli, E., & di Pellegrino, G. (2011). The neurobiology of intertemporal choice: Insight from imaging and lesion studies. *Rev. Neurosci*., 22(5), 565–574.

Shapiro, J. M. (2005). Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of Public Economics,* 89, 303–325.

Shiffrin, R., & Schneider, W. (1977). Controlled and automatic human information processing. *Psychological Review*, 84(2), 127–190.

Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., & Damasio A. (2005). Investment behavior and the negative side of emotion. *Pyschological Science*, 16(6), 435-439.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3–22.

Slovic, P. (Ed.) (2001). *Smoking: Risk, perception, & policy*. Thousand Oaks, CA: Sage Publications.

Smith, A. (1759) [1981]. The Theory of Moral Sentiments. In D. D. Raphael and A. L. Macfie (Eds.) Indianapolis, IN: Liberty Fund.

Spears, D. (2011). Economic decision-making in poverty depletes behavioral control. *The B. E. Journal of Economic Analysis & Policy*, 11(1), Article 72.

Strotz, R. H. (1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies.* 23, 165–80.

Sunstein, C. R., & Thaler, R. H. (2008). N*udge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.

Thaler, R. H., & Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy,* 89, 392–406.

Thornton, R. L. (2008). The demand for, and impact of, learning HIV status. *American Economic Review*, 98(5), 1829–1863.

Vervloet, M., Linn, A. J., van Weert, J. C. M., de Bakker, D. H., Bouvy, M. L., van Dijk, L. (2012). The effectiveness of interventions using electronic reminders to improve adherence to chronic medication: A systematic review of the literature. *Journal of the American Medical Informatics Association*, 19, 696–704.

Viscusi, W. K. (1992). *Smoking: Making the risky decision*. New York, NY: Oxford University Press.

Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5), 806–820.

Weinstein, N. D. (1987). Unrealistic optimism about susceptibility to health problems: Conclusions from a community-wide sample. *Journal of Behavioral Medicine*, 10(5), 481–500.

Wertenbroch, K. (1998). Consumption self-control via purchase quantity rationing of virtue and vice. *Marketing Science,* 17, 317–37.

White, J. S., Dow, W. H., & Rungruanghiranya, S. (2013). Commitment contracts and team incentives: A randomized controlled trial for smoking cessation in Thailand. *American Journal of Preventive Medicine*, 45(5), 533–542.

Wilde, P. E., & Ranney, C. K. (2000). "The monthly food stamp cycle: Shopping frequency and food intake decisions in an endogenous switching regression framework." *American Journal of Agricultural Economics*, 82(1), 200–213.

Zhou, X., Nonnemaker, J., Sherrill, B., Gilsenan, A. W., Coste, F., & West, R. (2009). Attempts to quit smoking and relapse: Factors associated with success or failure from the ATTEMPT cohort study. *Addictive Behaviors*, 34, 365–373.