Healthy Habits: Some Thoughts on The Role Of Public Policy in Healthful Eating and Exercise Under Limited Rationality

To appear as Chapter in book tentatively titled *Essays in Behavioral Public Policy*, edited by Adam Oliver, published by Oxford University Press.

November, 2011

Matthew Rabin

University of California—Berkeley

"Nothing so needs reforming as other people's habits."

— Mark Twain

"The diminutive chain of habit is scarcely heavy enough to be felt till it is too strong to be broken."

— simplification of Samuel Johnson quote used in 19th century temperance literature

"The child is the Father of the Man."

— William Wordsworth

"The second half of a man's life is made up of nothing but the habits he has acquired during the first half."

— Feodor Dostoevsky

"Men's natures are alike; it is their habits that separate them."

— Confucius

"My problem lies in reconciling my gross habits with my net income."

— Errol Flynn

Acknowledgments: I thank Alex Voorhoeve and seminar participants at the LSE for useful comments, and Mike Urbancic for valuable research assistance.

1 Introduction

In this chapter, I explore some possible policy implications concerning habitual activities relating to health—such as eating and exercise—that have (for good reason) become the subject of social and policy debate. I do so from the perspective of economic theory, empirical evidence, and with a focus on the implications of some recent research in behavioral economics.

I emphasize several themes. First, I outline a simple economic perspective on habitual behavior. Although it also accords with common sense, reasonable psychology, and empirical evidence, there is some tendency by social scientists and policymakers to neglect.

Second, as somebody who doubts the full rationality of many forms of habitual behavior, I discuss two basic errors that may lead people to engage in too many bad habits and too few good habits. One is widely discussed in psychology, behavioral economics, and the popular imagination: the human propensity to over pursue immediate gratification. Those forming habits simply care too much about current pleasures compared to what they themselves want in the long run. The other error has received less attention, and is less appreciated by behavioral economists and others, yet is likely to be crucial in the context of habit formation: people tend to underestimate (or at least under-attend to) the habitual nature of many activities. There is good psychology to suggest this is so. Furthermore, there is some evidence, and a very important logic, demonstrating that this underappreciation of the power of habits may be a tremendously costly error that policy makers may want to consider.

Third, I emphasize the extreme lack evidence on the degree to which eating and exercise are habit-forming, and on the degree to which people may behave irrationally in the face of this habit formation. Although I am not (anywhere remotely close to) an expert in empirical research in this area, I believe it is fair to say that there is remarkably little evidence in how

important the habits formed by children, adolescents, and young adults are in determining later behavior and preferences. It seems obvious, and almost surely true to some extent, that the role of habits is important. But it seems immensely important to know just how important. Even more importantly, do we have reason to believe that mature individuals and children and their guardians, fail to fully take into account the role of habits in their lives? Here too I think the psychological theory and lay impressions are on target. But whether people systematically err in their acquisition of any particular bad or good habit is rarely carefully established, even by those proposing policies that would seem predicated on the existence of such an error.

Finally, of course, if we believe that it is likely that many people are making these errors, and that these errors are both inherently unlikely to be meliorated by market incentives and are conducive to being addressed by policy, then we might want to study or implement policy interventions. I briefly discuss examples of possible interventions based on errors posited.

In the next section I discuss the most straightforward economist's conceptualization of what it means for activities to be habit forming, doing so within the utility framework. It is also my view of the most useful way to think about it. But while rational-choice economics assumes people behave 100% rational in the face of habit formation, in Section 3 I outline two specific types of errors (self-control problems and misprediction of preferences) that seem likely to lead people to over-indulge in bad habits, and in Section 4 I discuss more directly just how these errors affect habitual behaviors. In Section 5 I discuss my thoughts on the policy implications we might draw, and conclude in Section 6 with a brief review of empirical evidence—which is essentially a call for economists and policymakers to gather much more evidence.

My policy perspective in this article is that we should want to intervene in people's eating, exercising, and other behaviors insofar as it helps individuals achieve their own goals. This

means I will not discuss either the social benefits of healthful behavior, nor be judgmental about others' tastes. Although it is presumably better for society at large if people are healthier, I leave that public-policy goal as outside the scope of this paper.[1] My focus is on whether there is self-harm from activities. I also am uninterested in deciding whether some behaviors are good or bad based on my tastes or aesthetics. Mark Twain's quote above that "Nothing so needs reforming as other people's habits" strikes a cord as a permanently relevant ironic comment on policy debates in the domains we are considering. From the point of view of this essay, it is important that we find out if people are fatter than they want to be, exercise less than they want, smoke or drink more than they want, etc.—not what "we" want. Indeed, the two errors I review are "Twain-proof": "we" are the same as "they" for these errors. It is not my assumption that only readers who avoid the habits discussed will deem them worthy of regulating, but those of us who succumb to them as well.[2]

2 Preferences and Habits

How might we think about habitual behavior? As an economist who believes in the profound insight of utility theory for disciplined theoretical, empirical, and policy work on this topic, I will articulate a natural way for economists (behavioral or not) to think about habitual behavior. As argued below, positing the usefulness of modeling behavior in terms of people

[1] Economists are famously willing to make the point that bad health need not be costly to society. Indeed, it is my understanding that the current best estimate is that society benefits *financially* from tobacco-induced lung cancer. Speaking like an economist: if it kills people relatively cheaply and with relatively good timing—late in people's working lives but early in their social-security and medicare collecting lives (in the U.S. case)—it can be beneficial. Even if one buys into this calculus, I am not sure of the state-of-the-art estimates for obesity, but suspect they are likely to be different.

[2] This is in contrast to other types of errors policymakers and economists might study (such as over-trading in stocks) where we might be confident that "we" don't make the same errors as the population at large.

having well-defined "utility" for different outcomes which they are more or less trying to optimize is a very separate thing from saying that people *succeed* in that optimization. My perspective is, in fact, that the formal language of "preferences" is an efficient way to understand errors: without being clear about what people's goals are, it is hard to be clear about when it is right to say that somebody has made an error.[3]

The most basic notion of what it might mean for a good or an activity to be habit-forming, as modeled over the years, most famously in Becker and Murphy's (1988) notion of addiction, is that of intertemporal complementarities in the utility from consumption: an activity is habit-forming if the *marginal* utility a person gets today from consuming the good today is higher when she has consumed more of it in the past. Habits can be either "good", where the more you do of the activity the happier you are in the future, or "bad"—the more you do it, the less happy you are in the future. A good habit like exercise exhibits this shift—even if a certain exercise level starts out very unpleasant or even impossible (if you are out of shape, you may not be able to walk to work or climb 4 floors on the stairs), the more you do it the less unpleasant it becomes; it may even become a pleasant thing, or a compulsion.

Most of the focus of policy, of course, is on bad habits. The key insight of the utility

[3]In contrast to the utility-based notion of habit explored in this chapter, some psychological, neurobiological and economics approaches view habitual behaviors more as "unthinking", and not subject to optimization at all. Absent-mindedly reaching for cashews from a bowl on a coffee table, or lighting up a cigarette with friends at a bar without even thinking about it, can certainly in common parlance be considered a habit. An "addict" in some domain may be somebody who reacts to a cue by not optimizing in any sense. Indeed, research in both psychology (such as Baumeister (1994)) and economics (see Bernheim and Rangel (2004)) have emphasized some of the sundry ways that certain addictive or habitual behavior can be seen as variants of loss of volition. These issues are of course very real and potentially very important. But they concern basically understanding the daily behavior of those already addicted. They do not address the "Dostoevskian" point of asking how changes in the current environment may change what becomes habitual in the future, which is a major theme I emphasize below.

approach is that the notion that the marginal utility is higher when a person has consumed more in the past is perfectly consistent with the possibility that past consumption lowers utility *levels*. The essence of harmful habits is that current consumption decreases your future well being while causing you to desire those products more in the future. O'Donoghue and Rabin (1999b, 2000b) lay out a simplified, binary-choice form of the Becker, Grossman, and Murphy (1991) model of tobacco addiction, which applies equally to other types of habit-forming activities. Without presenting the mathematical model here, I can outline the basic features. We assume that each day a person can either take a "hit" or not take a hit. Our model incorporates two crucial components of harmful addictive activities. First, they involve *internalities*. The more common case is of *negative internalities*, where current consumption negatively affects future well being.[4] Negative internalities from consuming an addictive product may include future health, career, and personal problems, as well as "tolerance"—the fact that current consumption of a product lowers the the pleasure from future consumption. Habit-forming activities can also generate *positive internalities* if current consumption increases the instantaneous utility from future consumption. Second, harmful addictive activities involve *habit formation*: current consumption increases the marginal utility from future consumption. In other words, past consumption is assumed to increase the marginal instantaneous utility from current consumption—e.g., smoking cigarettes at age 16 increases the *marginal* utility from smoking a cigarette at age 17. The combination of negative internalities and habit formation means that as a person consumes more and more of an addictive product, she gets less and less pleasure from this consumption, yet she may continue to consume the product because refraining becomes more and more painful.

[4]Herrnstein and Prelec (1992) developed the term "internalities" in a related context to connote (by analogy with the economics term "externality") that consumption of a product now may impose benefits and costs on future selves.

Internalities come in many forms. Negative internalities include health problems due to over-eating or over-smoking. A person consuming cocaine or other drugs often exhibits "tolerance" in the sense that she derives less pleasure from a given level of consumption the more she has consumed in the past. Positive internalities arise from learning and other "investment goods", which we consume both for immediate gratification and for long-term benefits. Many cultural activities are presumed to have this property.

Activities can generate internalities without being habit-forming: eating cheesecake may generate negative internalities, and going to a museum may generate positive internalities, but neither is necessarily habit-forming. This is the big empirical puzzle that I think is so important for eating and exercise: although we are confident that tobacco, alcohol, and other substances are very significantly habit-forming, just how habit-forming eating and exercise patterns are is less clear.

What do these preferences imply about behavior? The primary assumption that economists have historically made in the context of habit formation is their standard and universally applied assumption: that people are fully rational. Before turning to potential errors people make in the next section, I first review what rationality implies about how people behave in the face of habit formation. A key prediction is that people are very much attuned to the future implications of their current behavior. This implies that people do not engage in activities that predictably lead to damaging bad habits whose pursuit exceeds the current pleasure of the activity.

Rather than directly address the question of whether people will rationally acquire bad habits, the original theory in Becker and Murphy (1988) paper famously focuses on a "steady state" analysis—looking for positive consumption levels such that a person would choose to *remain* addicted. In their empirical work, Becker, Grossman, and Murphy (1991, 1994) assume exogenous shocks in order to get people to become addicted. Although called a

rational-choice model of addiction, this rather non-standard focus in describing a behavior whose central premise is intertemporal tradeoffs and whose initial condition is always the same—non-addiction—might reasonably be considered a profoundly weak test of the rationality assumption. For those who think the big mistake in tobacco or alcohol addiction is to start the addiction (withdrawing from tobacco is deeply unpleasant; withdrawing from alcohol can kill you), the focus solely on "steady-state" levels is as if one explored a rational-choice model of suicide by asking solely whether it is rational for people to stay dead once they kill themselves. Other recent researchers—see Orphanides and Zervos (1995) and Wang (2007)—followed up Becker, Grossman, and Murphy's original research by addressing the more classical dynamic decision of whether to *become* addicted.

The early tests of whether addiction might be rational studied a particular prediction that is consistent with fully rational behavior, but not consistent with fully myopic behavior: that people respond to predictable changes in future prices, such as announced future tax increases. Becker, Grossman, and Murphy (1991, 1994) confirm Becker and Murphy's (1988) prediction of such future price effects. Although considered somewhat surprising that such future price elasticities were identifiable in the data—and despite skepticism that people would really know enough about future prices to react to them even if they wanted—Gruber and Kőszegi (2001) and Levy (2009) in fact support the finding of such future-price effects.

But as Gruber and Kőszegi (2001) also point out, the existence of future price effects suggests a rejection of *complete* myopia in the face of habit formation—demonstrating that people don't engage in habit-forming behavior as if their current smoking, injecting, or eating behavior had no effects whatsoever on their future attitudes towards smoking, injecting, or eating. But the fact that people are aware of such effects is predicted not only by complete rationality—it is a prediction of virtually all psychologically plausible theories of people's errors. Virtually any theory proposed by behavioral economists to modify the 100% rationality

model have also predicted such effects. And as theoretical and empirical research suggests—outlined in Gruber and Kőszegi (2004)—the amount of irrationality consistent with the observable future price effects would lead to wildly different policy and welfare conclusions. This is an absolutely central point: non-withstanding paper titles, and notwithstanding the default assumption of economics, there are no claims I am aware of in any papers that argue that existing evidence supports the policy conclusions of a 100% rational-choice model except if one a priori prefers 100% as a maintained empirical hypothesis, or one prefers it as the appropriate basis for policy when evidence is not sufficient. Becker, Grossman, and Murphy's (1991, 1994) estimates (as they themselves note in a footnote) yield future-price elasticities that are small enough to be suggestive that the fully rational model is not driving behavior. Gruber and Kőszegi (2004) show that the range of assumptions about self-control problems consistent with theirs and Becker, Grossman, and Murphy's evidence on tobacco includes optimal taxes that are much, much higher than what one would conclude under full rationality, and significantly higher than existing taxes. Levy (2009) estimates significant irrationality of both forms discussed below in tobacco consumptions, suggesting very different responses than if one were to assume full rationality.

In light of these findings, I now turn to a review of the two types of errors that theoretically may matter a lot in the context of good and bad habits. I first discuss the errors broadly, and then explore their implications more directly for habit formation.

3 Two Potential Errors

The classical economics model assumes that intertemporal preferences are *time-consistent*: A person's relative preference for well-being at an earlier date over a later date is the same no matter when she is asked. But casual observation, introspection, and psychological research all overwhelmingly suggest that preferences are often *time-inconsistent*: we tend to pursue

immediate gratification in a way that our "long-run selves" do not appreciate. The most important manifestation of time inconsistency is where greater weight is placed on immediate outcomes. O'Donoghue and Rabin (1999) label this latter tendency *present bias*.

Relative to time-consistent preferences, a person with present bias always gives extra weight to well-being *now* over any future moment, but discounts all future moments by the same discount factor. These preferences imply that at any moment a person may pursue immediate gratification more than she would have preferred if asked at any previous moment. That is, the person has a self-control problem. For example, when presented a choice between doing 7 hours of an unpleasant task on April 1 versus 8 hours on April 15, if asked on February 1 virtually everyone would prefer the 7 hours on April 1. But come April 1, given the same choice, most of us are apt to put off the work until April 15. Irrespective of the specific choice, time consistency requires that a person make the same choice on February 1 and April 1.

To examine intertemporal choice given time-inconsistent preferences, researchers have converged on a simple modeling strategy: for each point in time, a person is modeled as a separate "agent" who chooses her current behavior to maximize her current preferences, predicting how her future selves will behave.[5] In such a framework, an important issue arises: what does a person believe about how her future selves will behave? The answer to this question depends, of course, on beliefs about future selves' preferences. Two extreme assumptions have appeared in the literature to deal with the issue of beliefs about future behavior. *Sophisticated* people are fully aware of their future self-control problems, and therefore know exactly how their future selves will behave. *Naive* people are fully *un*aware

[5] To see the mathematical models used by economists to study present bias, including in the context of habit formation, see for example Laibson (1997), O'Donoghue and Rabin (1999a, 1999b, 2000a, 2000b), Gruber and Koszegi (2001), and Levy (2010).

of their future self-control problems and therefore believe their future selves will behave exactly as they currently would like them to behave.

Analysis of self-control problems typically involves comparisons of each of these types of people with the same long-run preferences: *time-consistent agents* are fully rational, as economists habitually assume. By systematically comparing these types of people, we are able to add insight into the role of self-control problems in habit formation, and to delineate how predictions depend on both self-control problems *per se* and on assumptions about foresight.

Although less integrated into formal economic models than self-control problems, recent research in economics has begun to focus on a second type of error that has been proposed (under different guises) by psychologists—and is of manifest relevance to habit formation. Do people rationally predict their own future tastes, and changes in those tastes? The wedge between our perceptions of what will make us happy and what does do so, and the wedge between what we think we will do and what we actually will do, are awfully big economic and social topics. Do we realize how happy more income and more work will make us? Do we have any idea what careers make us happy? Do we want kids? Are we going to get addicted if we start experimenting with this drug? Will we enjoy it? Do we fully predict and optimize with respect to the power of good and bad habits, such as exercise and eating habits?

Loewenstein, O'Donoghue, and Rabin (2003) argue that much of existing evidence can be summarized as a very general type of bias in prediction, which they formalize as *projection bias*: people tend to project their current tastes into the future, even in situations where those tastes change in predictable ways. In the context of habit formation, this is likely to have two very important implications: people may simply not attend to the way that current behavior forms habits, and (more subtly, and less central to my arguments here) people may

over-react to temporary fluctuations in feelings when planning future behavior.

To see how researchers have studied misprediction across domains, think of two general categories of ways that tastes change over time: temporary fluctuations, such as cravings and random moods, and longer-term systematic changes, such as addiction, health changes, and lifestyles and habits. The contention of researchers is that people underappreciate changes in their preferences, and hence falsely project their current preferences over consumption onto their future preferences.

Several studies, for instance, lend support to the folk wisdom that shopping on an empty stomach leads people to buy too much.[6] This phenomenon can be interpreted as a manifestation of projection bias: people who are hungry act as if their future taste for food will reflect such hunger. Read and van Leeuwen (1998) provide beautiful evidence of both "present bias" and "projection bias" in the context of hunger. Office workers were asked to choose between healthy snacks and unhealthy snacks that they would receive in one week, either at a time when they should expect to be hungry (late in the afternoon) or satiated (immediately after lunch).[7] Half of each of these two groups was approached either immediately after lunch, while the other half of each group was approached late in the afternoon.

Percentage of Subjects Choosing Unhealthy Snack:

		Future Hunger	
		Hungry	Satiated
Current Hunger	Hungry	78%	56%
Hunger	Satiated	42%	26%

The above chart indicates choices that subjects made a week in advance, and clearly indicates that people were influenced by their current hunger as much or more than their

[6] See, e.g., Nisbett and Kanouse (1969) and Gilbert, Gill and Wilson (1998).

[7] The healthy snacks were apples and bananas, the unhealthy snacks were crisps, borrelnoten, Mars Bars, and Snickers Bars.

completely predictable future hunger state. Notice that the findings of influence of current hunger state has nothing to do *per se* with self-control problems—whether they "should" or should not want healthy foods (which are the labels the authors used, but of course did not present to subjects—who were just asked to choose among a list of items). The point is that for the choice about the exact same future situation, subjects were influenced by their current hunger. This is a surprising result given how familiar these office workers presumably are—like the 7 billion other people—with time-of-day fluctuations in hunger.

The researchers however did collect evidence suggesting self-control problems and time inconsistency with respect to eating: unbeknownst to the office workers, when the snacks were brought back next week, they were allowed to change their minds. Of those who chose healthy snacks ahead of time, 44% chose to switch to unhealthy items. Of those who chose unhealthy ones, only 3% wanted to switch to healthy! This, in a nutshell, indicates the time inconsistency. Prospectively people wanted to eat healthier than they wanted to eat later—at the moment of truth.

A fascinating study by Giordano, Bickel, Lowenstein, Jacobs, Marsch, and Badger (2002) (also analyzed in Badger, Bickel, Giordano, Jacobs, Loewenstein, and Marsch (2007)) similarly finds an indication of both present bias and projection bias. They studied 13 long-time adult heroin addicts who had been regularly receiving BUP—a medication that reduces craving for heroin and is used in aiding with heroin withdrawal. Addicts currently deprived of BUP like it more than those not currently deprived. Over an 8-week period, each subject was asked whether s/he would prefer each of 12 different amounts of money (ranging from \$0 to \$100) to a second dose after receiving an initial dose. (This second dose is still attractive to addicts.) Subjects were told (truthfully) that one of their choices, randomly selected, would be implemented. Hence, they had the incentive to choose according to their true preferences.

Half the time subjects were asked when "deprived"—2 hours before receiving their sched-

uled dose for the day. In the "satiated" condition, they were asked right after receiving their dose. Half the time subjects were asked for their willingness to pay for the additional dose today (but, again, in both cases after the initial dose, whether they had gotten that initial dose yet), and other half asked for getting it 5 days hence. In all conditions, the "state" of craving for the dose they'd receive were identical, but the timing and the craving at the time of decision differed. They found an average willingness to pay, as a function of both their state when asked and the imminence of delivery of the dose as follows:

		When they would get the dose	
		Today	5 Days from now
Current	Deprived	\$75	\$60
Craving	Satiated	\$50	\$35

These results too indicate a deep neglect by even very experienced addicts of even very predictable changes in future craving levels.

Along with this projection bias, these results also indicate present bias: the fact that, irrespective of the craving state at the time asked, these addicts were willing to pay \$15 more for a second dose on the same day than for a week later clearly indicates a taste for immediate gratification.

Unfortunately, it is harder to find convincing evidence of longer-term habit formation, and especially to find convincing evidence of *misprediction* of such longer-term changes. I discuss the limited evidence in the context of habit formation below.

Finally, before moving on to studying the implications of these errors, I emphasize something related to the dictum "De Gustibus Non Est Disputandum": both of these errors discussed above are errors in implementing the preferences people have, not "wrong preferences". These models take as given preferences that people have, and predict behavior and errors given those.

4 Implications of These Errors for Habit Formation

Even fully rational people can form "bad habits", and unless we want to moralistically embrace Mark Twain's observation, but miss the irony, our interest as a society ought be focused in a large part on understanding *mistaken* bad habits. Yet there are several reasons to believe that many bad habits—and missing good habits—are due to the two errors discussed above, as well as other errors. First, at a casual empirical level, people have many bad habits that they vociferously say they don't want to have. This does not prove anything—people don't always say in words what they truly want, and they may not actually know.[8] But it would be policy folly and scientific silliness to ignore predictable expressions of people wishing things were different for them. The fact that parents and other loved ones wish people had better habits also means something.[9] And more direct measures of well-being, from quantitative evidence like Gruber and Mullainathan's (2005) arguments that people who are heavily taxed on tobacco are on average happier, to your own observations about the seeming happiness of people you know, indicate people may not have the levels of habits that make them happiest. Finally, the fact that even the most dedicatedly rational-choice research by economists studying habits falls so incredibly short of being convincing that pure rationality rules the day, should be a strong signal.

What might the two types of errors discussed above say about habit formation? First, the direct implication of having a self-control problem is a tendency to over-consume goods and activities with negative internalities, because of underweighting the future health costs

[8] And the fact that people say honestly that they wish they were thinner, or were in better physical shape, isn't per se the question we are interested in. Presumably everybody wishes they were healthier. We need to know whether people wish they paid the cost to achieve these goals, not just that they wish the goals were achieved without cost.

[9] But it is likewise not proof; there are certainly wishes parents have for children that are themselves mistakes, or self-serving.

relative to the current pleasures. People may eat too much, walk too little, etc. This simple and direct fact matters a lot, and is presumably what many people have in mind when they consider over-consumption to be a problem of societal and policy concern. This overconsumption inspires some of the policies I discuss below.

But it is not about *habitual* nature of eating and exercise per se. This tendency to over consume would be true even if these activities were not habit-forming. Importantly, the case that *habit formation* itself exacerbates the effects of self-control problems on overconsumption is not at all straight forward. If people are fully aware that some activity is habit-forming, then short of complete (and unrealistic) unconcern about their future, it will change their behavior. Even a person tempted by cake or by driving rather than walking who knows that current behavior will influence her future behavior and outcomes will take those effects into account. The observation that the misbehavior of over-eating while young, say, is more costly if it leads to a taste for excessive eating while old, is correct as far as it goes. But that does not imply that the case for societal attention to correcting these self-control problems is any greater by dint of its habit formation. If young people (or their guardians) realize the stakes are larger, they themselves will reduce the levels of an activity.

Consider a teenager who rightly or wrongly believes that a drug is harmful to her health and future, but *not* addictive. She then perceives the cost of using the drug as the direct costs, which she may count less than she should because of self-control problems. But if she believes that a drug is addictive, she'll now also recognize the cost of using the drug—that it will lead her to crave the drug in the future. This may in fact deter her without outside intervention. She will be *less* likely to engage in addictive activities than equally pleasurable non-addictive activities.

There is a similar—more subtle, but in fact more dramatic—reversal of common intuition regarding the effects of naivety about self-control. With any activity that is not habit

forming, but involves trading off current pleasures against future health and other costs, then—in the absence of commitment devices—naivety has roughly no effect. Essentially, if eating desert today has roughly no effect on either your taste for desert in the future or on how bad future desert will be for your health, then whether you are aware of your propensity to eat desert tomorrow will not influence your tendency to eat desert tonight. But if behavior is habit forming, your theory of how you will respond in the future to habits you acquire today, will influence your perceived cost of indulging today.

What effect will it have? Here, the common intuition that naivety about self-control is especially damning is not necessarily right. In fact, in the most basic settings, it is backwards. One might think that if you naively think you will avoid developing a bad habit in the future than you will be insufficiently scared about acquiring the habit. But it is more painful to break yourself of a habit than to never acquire one; this is more or less what we mean by bad habits. If you naively believe that you will have the willpower to quit even if you develop a habit, then you perceive the cost of acquiring the habit as *more* costly than if you think you will continue to indulge anyhow. Likewise for good habits: letting yourself get out of shape is *more* tempting if you sophisticatedly realize you will just get out of shape in the future anyhow than if you naively think you'll make the difficult reinvestment in getting back in shape.

In many simple situations, then, naivety about future self-control may help you rather than hurt you in avoiding bad habits. In other situations, sophistication may help you. If you have more complicated tastes—especially if you realize that your tastes will change with age and you will stop being tempted—then sophistication may help you avoid naively indulging in youth when it is attractive but then building an irreversible bad habit. And, related to this point, it turns out that sophistication may help tremendously in reforming a habit that you (for some reason) acquire. Intuitively, naivety can lead you to persistently,

over-optimistically predict that you will "quit tomorrow"—making you less motivated to quit today. The decision when to start withdrawing from a habit that you plan to withdraw from is psychologically and mathematically very similar to the type of "procrastination" discussed in Akerlof (1991) and analyzed in detail in O'Donoghue and Rabin (1999). In situations where you must do something once involving an up-front cost with long-term benefits—precisely the case of changing a bad habit—naivety about your self-control problem almost always hurts you by making you repeatedly procrastinate under the false hypothesis that you will successfully complete the task tomorrow.

In sum, neither self-control problems per se nor naivety about future self-control problems necessarily lead to special mistakes in light of habit formation. People over-indulge in unhealthful behaviors because of self-control problems. But it is primarily the unhealthfulness per se—not the habit component—that causes the problem. Self-control problems predict that people overconsume unhealthy goods and activities, but by themselves there is no simple prediction that overconsumption is worse for habit-forming goods than for non-habit-forming goods.

The implications of projection bias, however, is less ambiguous: it does not lead to over-indulgence in non-habit-forming vices, but *does* lead people to engage in too many bad habits. Although less sharply focused on by both researchers and policymakers, I believe it is likely to be the larger error—although certainly there are cases where both in combination may do much more damage than either mistake alone would. Simply put: some evidence, and lots of intuition, suggest that people underpredict—or, if they intellectually predict, are too inattentive to it—the effects of habits. If you don't imagine at the time of choice how much your tastes will change if you change your lifestyle—if you don't believe that you'll ever come to enjoy veggies, or be satisfied with less to eat, or that exercise at the gym or walking up that flight of stairs will become easier over time—then you will under-invest in

good habits and acquire many unplanned bad habits.

An interesting ancillary problem caused by projection bias for bad habits seems of clearest relevance to dieting: insofar as people underappreciate variation in hunger levels—as hinted by the Read and van Leeuwen study cited above—then you can not only overeat but "over diet" in an ineffective cycle of dieting and binge eating. When you just (over)ate for the day, and are very full, you may naively imagine that it is easy to refrain from eating. If you can't eat another bite right now, you may feel as if you can never eat another bite! You start the diet. And then you get hungry. Not only does the hunger lead you to eat (as it should), but you feel at such moments that it is unimaginable to go long without eating much—you may give up altogether on the diet. And so you eat until you can't eat another bite ...

5 Policy Implications of Those Errors

I reiterate two points from earlier. First, unless we believe that it is likely that people are making mistakes, the fact that some activities are habit forming does not (in any way that I can understand) heighten the case for policy intervention, regulation, or paternalism.

Second, if we decide we'd like to deter some activity, we should never forget the power of prices. The most practical policy we may employ if we reach the conclusion that people are doing too much of bad habits or too little of good habits: tax or otherwise deter the bad habits, and subsidize or otherwise support the good habits. If we want to get people to do less unhealthy eating, we should make it more costly; if we want people to do more exercise, we ought to make it cheaper. Of all those quoted above, Errol Flynn is classically considered the least wise, and his observation that "My problem lies in reconciling my gross habits with my net income" does not satisfy our sense of depth in understanding habits. But he wisely noted that not having enough money to indulge his habits was a problem and surely recognized it is more of a problem for expensive habits than cheap ones. Flynn may

have recognized something deep that the others did not: making something more costly is a pretty robust way to get people to do less of it.

Quite how to go about making bad habits more costly is of course harder, and certainly old-fashioned taxation and subsidies aren't the only instruments we ought to think about. But I worry that many concerned people might outsmart themselves while looking for more sophisticated determinants of behavior and ignore the tried and true method of literally or metaphorically using prices to influence behavior. Taxes and subsidies work to influence behavior in virtually any domain and irrespective of whether and how people are irrational.

What does the nature of habit formation, and the potential errors exacerbated by habit formation, suggest about when and how to impose the prices? There is a simple logic of how to intervene that many of us have advocated that would work in principle for bad habits, but is much harder to see how to implement with either eating or exercise (or alcohol) than with more sharply addictive activities like smoking. We are likely to get more bang for the buck by targeting people when they are younger than when they are older. If we want to stop bad habits with as little effort and as little invasiveness as possible, we should stop them before they start. Indeed—and very importantly—for the various paternalistic and non-paternalistic goals, it is eminently plausible that—although we should want to stop bad habits and addictions before they start—we should not necessarily try to shut them down after they start. For good habits, the principle is that we should probably, when feasible, try to incentivize their creation but not necessarily subsidize their continuation.

The structure of habits and how to affect them with taxation and other regulations is an especially important issue for another important goal of public policy: progressive taxation and redistribution. If our ancillary goal is to change behavior without too much undesirable redistributional effects, some argue (often sincerely, sometimes not so sincerely) that they do not want to tax vices because (as is empirically so) many of these vices are more prevalent

among lower-income people. If we tax fatty food or tobacco, we are, in many more developed countries, disproportionately taxing the poor. As Gruber and Kőszegi (2004) argue in the context of tobacco, however, there is an incompleteness to the intuition people have that the monetary tax incidence may be regressive, because poorer people are more likely to smoke or eat fatty food than wealthy people. The financial cost, in fact, may get things backwards. Simply put: a smoker pays a huge amount in taxes over the course of his lifetime. Raising taxes may get young people never to take up tobacco, and save them decades of extra taxes. The net tax burden from raising taxes on bad habits may be negative, and therefore be a means of redistributing money *towards* the groups engaging most in those bad habits.

Of course, there is a clear cost to taxing people, especially if many people are consuming goods at the right level. If we impose a heavy tax on crisps, then the many people who enjoy crisps will have to pay more. Economists have good reasons to argue that when people are 100% rational, we should (depending on certain details of the structure of demand functions) tax different commodities at the same rate. If we would like to raise revenues corresponding to an average VAT of 10%, we often want to do so by having a uniform VAT. If people were 100% rational, we would not, for instance, want to have a tax of 9% on most items and 200% on unhealthy food if that raised the same amount of taxes.

Here there is an issue with the way economists are and are not contributing to a mature and empirically grounded debate about the benefits and costs of "sin taxes". There are two interelated doubts I have about economists' contribution to this debate, or at least something to watch out for as the debate over such taxes sharpen. First, related to both Gruber and Kőszegi (2006) and O'Donoghue and Rabin (2006), adding even a bit of realistic self-control problems that are manifestly at least as consistent with observed demand as the 100% rational model dramatically changes what economic theory says is the optimal tax. That is, economists have only to be a little bit wrong in our extreme rationality assumptions

in order for our prescriptions for taxes to be pretty significantly wrong, and for the possibility that significant taxes on some unhealthy items are warranted. But this is closely related to a second problem with frequent framing by economists: even *assuming* that people are 100% rational, all researchers ought be forced to *quantify* the harm done by "distortionary" taxes. That is, suppose that 100% of those who ate junk food were 100% rational in doing so. If we doubled the price of such food via taxes and lowered taxes on other items to keep revenue the same, how bad would that be, using the same types of measures (like "consumer surplus") that economists already use? My sense is that the answer is often "not bad at all". Insisting that economists correctly identify what their own theories say is the *size* of the harm done from proposed policies I believe will change the nature of the debate, for both the societal and scientific good. Economists as a group have perhaps come up short of their usual scientific standards in their initial reactions to some of these debates. I think that is largely for the forgivable and inevitable reason that the potential benefits of many of the proposals are not something that economists have the familiarity and formal means to engage with. Hence they simply report what their theory says is the best policy, rather than the costs of what they consider pointless departures from that best policy. Some of it may be for the (maybe marginally less forgiveable, nor inevitable) reason that economists are so worried about supporting such policies for reasons that are unrelated to existing economic theories or evidence that they don't want to come clean about what those existing theories and evidence actually indicate. The two questions—the optimality of even fairly significant-seeming taxes for relatively small degrees of irrationality, and the fact that even maintaining 100% rationality does not imply a huge cost—are related: theory tells us that the welfare costs of "sin taxes" are small ("second order") to fully rational people and the welfare benefits are large ("first order") to those making the types of mistakes posited.

This, in turn, helps indicate a specific policy that will make sense in light of the errors

discussed above, and to highlight a major theme in "behavioral public policy" research. Under various names—cautious paternalism, conservative paternalism, benign paternalism, or libertarian paternalism, and "nudging"—research has proposed ideas of how to implement policies that do the least amount of harm to those who might be behaving fully rational and those behaving badly. To illustrate, let us abstract away from the very real implementation problems (black markets, arbitrage, the costs of ID checks, etc.), and consider the following thought experiment that could be turned into a real policy experiment—or a real policy. Without taxing crisps or soda or cigarettes any more in total than we already do, consider merely changing the life cycle of the tax. Instead of (say) 10% tax on unhealthy items for a person's entire life, consider heavy taxes for young people for these items, and no taxes when older, in a way calculated to leave the total tax burden the same overall if people do not change their behavior. What would happen, according to different theories of motivation? If young people are acting according to fully rational models, fully realizing the habits they are forming and the costs they are incurring, then they will be made no worse off. Indeed, there is a behavioral prediction of the rational model: they will either keep consuming a lot in their youth and in their adulthood just like they did before, or they will stop in their youth and then start in their adulthood. But either way, economic theory based on full rationality says they will be just as well off as before. How might people who have self-control problems or projection bias behave? The prediction is that they are very likely to decrease consumption dramatically both in their youth *and* thereafter. This is because the prediction of these alternative models is that those who were forming these habits when young (at least the ones who were close to indifferent before) were not planning to do so. If people don't realize they will develop a lifelong habit as strong as they will, then they never thought they were going to pay taxes later in life just because of early consumption. We've now changed taxes so that they will pay the same *actual* total price of the habitual behavior that they were in

fact going to pay under the old taxes, but perhaps not the taxes they intended. If their old behavior was based on not realizing they were starting a habit, this would tax them out of consumption. If their old behavior was based on correct realization of their habit they were forming, it won't stop them consuming—and won't hurt them. More feasible and finetuned cousins of this hypothetical policy would improve the well-being of those overconsuming bad habits because of misprediction and would not hurt significantly (or at all) those doing so rationally.

This proposed policy, if feasible, would also of course provide empirical evidence on the role and rationality of habit formation in behavior. But my focus on the case of young people leads to a second way I'd like to see economists more forcefully enter the policy debate on the role of habit formation in unhealthful behavior. Here I want to remove my behavioral-economist, theoretical, and welfare-analytic hats, to urge far-greater focus on a simple, yet important, basic fact: a huge amount of our behavior and welfare in the second half of our life is—in a coherent and important sense—determined by our choices in the first half of our life. Whether seen through William Wordsworth's famous line "The child is the Father of the Man," or Dostoevsky's plain empirical statement "The second half of a man's life is made up of nothing but the habits he has acquired during the first half," this is a simple empirical hypothesis that is of clear moral and policy relevance.

Is it true? I end this essay not with exploring more policy implications, but with where I think some of the longer-term policy implications ought to begin: further empirical evidence. Just how important *is* the role of habit in eating, exercise, and other healthful activities? Is it true that (say) changing children's eating habits will last them their whole life? How big is the effect? Again emphasizing my lack of familiarity with much of the empirical literature, I am still confident that very little convincing economic research exists trying to carefully observe the role of habit and prediction of habit in exercise, and almost none on long-run

eating habits.

6 Empirical Evidence on Habit Formation

There is one obvious approach we do *not* want to take to establish the role of habit in eating and exercise: compare the later behavior of those who eat healthy early in life to those who do not eat healthy early in life. These will almost surely be positively correlated, and one often hears statistics of this sort presented as if it is laden with meaning. It is true that kids who eat more crisps tend to turn into the adults that eat more crisps. But we can be positive that much of this is not causal: children who like crisps better turn into adults who like crisps better, and so their intrinsic taste alone guarantees their consumption. What we want to know, to understand the benefits of policy trying to influence early life healthfulness, is whether an "exogenous" shift in eating crisps of the sort induced by a change in policy will change late life crisp eating. To many of us inclined to believe in both the likelihood that healthful behavior is habitual, and that much of it is mistaken and hence likely to benefit from policy oriented towards correcting it, unconvincing claims that correlation between behaviors early in life and late in life implies a causal role of habits undermines the scientific case for whether and how to pursue corrective policies. Nor is the heuristic that "surely" some of the correlation is causal very convincing. In fact, we care a lot how much of it is. It is even possible that there is a positive correlation but a negative causal role. Children with a sweet tooth may grow into adults with a sweet tough, even if eating sweets when young makes them less keen on sweets when old.

This is a tremendously hard empirical problem to tell whether early-life behavior really turn into causal habits. In the domain of eating, it is especially difficult. One would want to know whether (say) two schools or similar school districts that had different school food policy *unrelated to the school populations* generated different habits later in life. This is a

typical, and typically hard, empirical-economic problem. Another common technique is to try to look at natural experiments of some policy change that we know or suspect has an immediate affect on behavior between two different close groups, and see the long-run effect. If some place banned soda machines from schools, for instance, then investigating whether years later there is a noticeable jump in habits of people at a certain age to those one year younger could indicate a habit-forming change in behavior at a formative age.

Another approach is to attempt to use the effect of past price changes on current behavior, along the lines of the tobacco literature. Earlier economic identification of habits often *assumed* rational expectations. That means essentially looking for future price effects: something is habit forming (for a completely rational and informed consumer) in proportion to how much the person will react to perceived future prices in changing current behavior. But the approach of looking at long-run effects of past prices does not assume rational expectations.[10] This is, empirically, a very hard thing to do. It is very difficult to separate out the effects of a lifetime of different prices, and it is very hard to find people who faced different prices for entirely "exogenous" reasons.

The existing literature on the addictive nature of smoking attempts to make progress in that domain. While acknowledging the challenges to confident identification, Gruber (2001a) probably gives the best guess. Using price variation across states and information on where people grew up, he estimates that something like 25-50% of exogenous increasing in youth smoking translates into adult smoking. That is, of every 100 young people priced out of smoking due to the range of tax changes he observes in the data, 25 (conservative estimate) or more likely 50 of them will never become smokers. Although I find Gruber's (2001a) estimates more convincing, see Glied (2003, 2004), who argues that most of the effects of price-based tobacco reduction among youth disappears by age 40. Although it relies on

[10] See Pollak (1970).

many more assumptions about the utility function and empirical estimation, Levy's (2009) estimates can be used to surmise that if we could stop people smoking at all before the age of 30, then two thirds of them would never begin. Levy (2009) also provides evidence of the mistakes involved in smoking among youth: inferring from the price responsiveness and other identification techniques, he suggests that beginning smokers suffer from projection bias, and are not predicting their future addiction much at all.

Cigarettes are, of course, both more habit forming and far easier to study concretely than "eating habits". It will be much harder to really learn about the long-run nature of eating or exercise habits. But to get a grip on shorter-run effects, some exciting new research by experimental economists has shed some light. Charness and Gneezy (2009) conducted a simple experiment at a gym that indicates that exercise is habit-forming. During a treatment period, they randomly paid some non-gym-goers money to go to the gym, while others in their experiment were treated similarly but not incentivized to go to the gym. Unsuprisingly (prices again!), those paid to go the gym were more likely to go during the incentives period. Yet, surprisingly or not, this group was also more likely to go even after the incentives were removed. Although this effect wore off over time, there seemed to be a clear habituation to exercise. Acland and Levy (2010) replicated these results in a very similar design, also showing habituation that diminished over time. But they also had subjects in their study make incentivized predictions about their own future behavior at various times and incentive conditions, and were able to show that people were both naive about their self-control problem, and (importantly) that they underappreciated the good habit formation associated with exercise. That is, before they were given incentives to go to they gym, they did not realize that this would make them more eager to go to the gym later. Although the time scales in these studies are necessarily not as long as we'd like, they provide evidence for habit formation and a guide to further research.

7 References

Acland, Daniel and Matthew R. Levy. (2010). "Habit Formation, Naiveté, and Projection Bias in Gym Attendance." Mimeo.

Ainslie, G. (1975). "Specious reward: A behavioral theory of impulsiveness and impulse control." *Psychological Bulletin* 82(4): 463-496.

Ainslie, G. and R. J. Herrnstein (1981). "Preference reversal and delayed reinforcement." *Animal Learning & Behavior* 9(4): 476-482.

Akerlof, George. (1991). "Procrastination and Obedience." *American Economic Review* 81(2): 1-19.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremey Tobacman, and Stephen Weinberg. (2001). "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15(3): 47-68.

Badger, Gary J., Bickel, Warren K., Giordano, Louis A., Jacobs, Eric A., Loewenstein, George, and Marsch, Lisa (2007), *Journal of Health Economics* 26(5), 865-876, September 2007.

Baumeister, Roy F., Todd F. Heatherton, and Dianne M. Tice. (1994) *Losing control: How and why people fail at self-regulation.* San Diego, CA: Academic Press.

Becker, G. S., M. Grossman and K. M. Murphy (1991). "Rational Addiction and the Effect of Price on Consumption." *American Economic Review* 81(2): 237-41.

Becker, G. S., M. Grossman and K. M. Murphy (1994). "An Empirical Analysis Of Cigarette Addiction." *American Economic Review* 84(3): 396-418.

Becker, G. S. and K. M. Murphy (1988). "A Theory of Rational Addiction." *Journal of Political Economy* 96(4): 675-700.

Bernheim, B. Douglas and Antonio Rangel. (2004). "Addiction and Cue-Triggered Decision Processes." *American Economic Review* 94(5): 1558-1590.

Charness, Gary and Uri Gneezy. "Incentives to Exercise." *Econometrica* 77(3), 2009: 909-931.

Confucius. (1995 [ca. 475 BCE - 221 BCE]). *The Analects.* New York: Oxford University Press, 1995, book 17, chapter 2.

Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang. (2007). "Projection Bias in Catalog Orders." *American Economic Review* 97(4): 1217-1249.

Delavan, Edward C., ed. (1884). *Temperance Essays, and Selections from Different Authors.* New York: The National Temperance Society and Publication House, p. 269.

Dostoevsky, Fyodor. (1995 [1872]). *Demons.* trans. by Richard Pevear and Larissa Volokhonsky. London: Vintage.

Gilbert, Daniel T., Michael J. Gill, and Timothy D. Wilson. (2002). "The Future Is Now: Temporal Correction in Affective Forecasting." *Organizational Behavior and Human Decision Processes* 88(1): 430-444.

Giordano, L., W. K Bickel, G. Lowenstein, E. A. Jacobs, L. Marsch, and G. J. Badger. (2002). "Mild opioid deprivation increases the degree that opioid-dependent outpatients discount delayed heroin and money." *Psychopharmocology* 163: 174-182.

Glied, S (2002), "Youth tobacco control: reconciling theory and empirical evidence," *Journal of Health Economics* 21(1) 117-135, January 2002.

Glied, S (2003), "Is Smoking Delayed Smoking Averted?" *American Journal of Public Health* 93(3), 412-416, March 2003.

Gruber, Jonathan. (2001a). "Youth Smoking in the 1990's: Why Did It Rise and What Are the Long-Run Implications?" *American Economic Review* 91(2): 85-90.

Gruber, Jonathan, ed. (2001b). *Risky Behavior Among Youths: An Economic Analysis.* Chicago: University of Chicago Press.

Gruber, Jonathan and Botond Kőszegi. (2001). "Is Addiction 'Rational'? Theory and

Evidence." *Quarterly Journal of Economics* 116(4): 1261-1303.

Gruber, Jonathan and Botond Kőszegi. (2004). "Tax incidence when individuals are time-consistent: the case of cigarette excise taxes." *Journal of Public Economics* 88(9-10): 1959-1987.

Gruber, Jonathan and Sendhil Mullainathan. (2005). "Do Cigarette Taxes Make Smokers Happier?" *Advances in Economic Analysis and Policy* 5(1): article 4.

Herrnstein, R. J. and D. Prelec (1992). "A Theory of Addiction," in *Choice over time*, G. Loewenstein and J. Elster, eds., Russell Sage Foundation, New York, NY, US: 331-360.

Johnson, Samuel. (1809). "The Vision of Theodore, the Hermit of Teneriffe," in *The works of Samuel Johnson, L.L.D., in twelve volumes.* Vol. 2. Boston: Hastings, Etheridge and Bliss, p. 455.

Laibson, David. (1997). "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112(2): 443-477.

Laibson, David. (1998). "Life-cycle consumption and hyperbolic discount functions." *European Economic Review* 42(3-5): 861-871.

Laibson, David, Andrea Repetto, and Jeremy Tobacman. (1998). "Self-Control and Saving for Retirement." *Brooking Papers on Economic Activity* 1: 91-196.

Levy, Matthew R. (2009). "An Empirical Analysis of Biases in Cigarette Addiction." Mimeo.

Loewenstein, G. Ted O-Donoghue, and Matthew Rabin. (2003). "Projection Bias In Predicting Future Utility." *Quarterly Journal of Economics* 118(4): 1209-1248.

Loewenstein, G. and D. Prelec (1992). "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics* 107(2): 573-597.

Nisbett, R. E. and D. E. Kanouse. (1969). "Obesity, Food Deprivation and Supermarket Shopping Behavior." *Journal of Personality and Social Psychology* 12(4): 281-294.

O'Donoghue, T. and M. Rabin. (1999a). "Doing It Now or Later." *American Economic*

Review 89(1): 103-124.

O'Donoghue, T. and M. Rabin. (1999b). "Addiction and Self Control." in *Addiction: Entries and Exits*, Jon Elster, editor, Russel Sage Foundation.

O'Donoghue, T. and M. Rabin (2000a). "The economics of immediate gratification." *Journal of Behavioral Decision Making* 13(2): 233-250.

O'Donoghue, T. and M. Rabin. (2000b). "Addiction and Present-Biased Preferences." Mimeo.

O'Donoghue, T. and M. Rabin. (2001a). "Risky Behavior Among Youths: Some Issues from Behavioral Economics." in *Risky Behavior Among Youths: An Economic Analysis.* Gruber, Jonathan, ed. Chicago: University of Chicago Press.

O'Donoghue, T. and M. Rabin. (2001b). "Self Awareness and Self Control." in *Now or Later: Economic and Psychological Perspectives on Intertemporal Choice*, Roy Baumeister, George Loewenstein, and Daniel Read (eds.), Russell Sage Foundation Press.

O'Donoghue, T. and M. Rabin (2003). "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes." *American Economic Review* 93(2), 2003: 186-191.

O'Donoghue, T. and M. Rabin. (2006). "Optimal Sin Taxes," *Journal of Public Economics* 90(10-11), November 2006: 1825-1849.

Orphanides, A and D. Zervos (1995). "Rational Addiction with Learning and Regret." *Journal of Political Economy* 103: 739-758.

Phelps, E. S. and R. A. Pollak. (1968). "On Seocnd-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35 (2): 185-199.

Pollak, R. A. (1968). "Consistent Planning." *Review of Economic Studies* 35(2): 201-208.

Pollak, R. A. (1970). "Habit Formation and Dynamic Demand Functions." *Journal of Political Economy* 78(4 July/August): 745-763.

Pollak, R. A. (1978). "Endogenous Tastes in Demand and Welfare Analysis." *American*

Economic Review 68(2 May): 374-379.

Read and van Leeuwen. (1998). "Predicting Hunger: The Effects of Appetite and Delay on Choice." *Organizational Behavior and Human Decision Processes* 76(2): 189-205.

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

Strotz, R. H. (1955). "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23(3): 165-180.

Sunstein, Cass R. and Richard H. Thaler. (2008). "Nudge: Improving Decisions About Health, Wealth, and Happiness." New Haven, CT: Yale University Press.

Twain, Mark. (1894). *Pudd'nhead Wilson.* New York: Charles L. Webster & Company.

Wang, R. (2007). "The Optimal Consumption and the Quitting of Harmful Addictive Goods," *The B.E. Journal of Economic Analysis and Policy* 7(1), article 15.

Wordsworth, William. (1807). "My Heart Leaps Up When I Behold," in *Poems, in Two Volumes.* London: Longman, Hurst, Rees, and Orme.