

NBER WORKING PAPER SERIES

INTERTEMPORAL CHOICE

Keith Marzilli Ericson
David Laibson

Working Paper 25358
<http://www.nber.org/papers/w25358>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2018

We are grateful to the volume editors Doug Bernheim and Stefano DellaVigna, who offered terrific insight and guidance on our chapter. We also received excellent advice from George Loewenstein and Matthew Rabin. We also thank all of the participants at the SITE 2016 Conference who shared feedback with us. Much of our thinking on this topic has emerged and evolved from our collaborations and conversations with Jonathan Cohen and John White. Lea Nagel provided outstanding research assistance. We gratefully acknowledge financial support from The Pershing Square Fund for Research on the Foundations of Human Behavior, the Boston University Angiola M. Noe Research Fund, and the National Institute of Aging R01 (R01AG021650). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Keith Marzilli Ericson and David Laibson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Intertemporal Choice
Keith Marzilli Ericson and David Laibson
NBER Working Paper No. 25358
December 2018
JEL No. C90,D14,D15,D60,D91

ABSTRACT

Intertemporal tradeoffs play a key role in many personal decisions and policy questions. We describe models of intertemporal choice, identify empirical regularities in choice, and pose new questions for research. The focus for intertemporal choice research is no longer whether the exponential discounted utility model is empirically accurate, but, instead, what models best explain the robust behavioral deviations we observe. We introduce the term “present-focused preferences” to describe the large class of models that prioritize present flows of experienced utility. Present-focused preferences need not coincide with a preference for commitment or dynamically inconsistent preferences. Present-bias is a special case of present-focused preferences.

Keith Marzilli Ericson
Boston University Questrom School of Business
595 Commonwealth Avenue
Boston, MA 02215
and NBER
kericson@bu.edu

David Laibson
Department of Economics
Littauer M-12
Harvard University
Cambridge, MA 02138
and NBER
dlaibson@gmail.com

1. Introduction

Most decisions have consequences that play out over time. How much should I spend today and how much should I save? How many hours should I work on a problem set tonight and what work should I postpone? Should I have a candid conversation with an under-performing co-worker or delay the awkward interaction? Is it worth getting out of bed to take my medicine, or is it OK to skip a night? Should I exercise this afternoon, or check all of my social media accounts today and exercise tomorrow?

Whether in the workplace, the marketplace, on vacation, or at home, almost all decisions have an intertemporal dimension. If one makes these decisions with any foresight at all, it is necessary to somehow weigh utility flows (i.e., costs and benefits) that occur at different points in time.

These questions also dominate many of our world's leading policy questions. How much to invest for the future is at the heart of myriad policy issues, including education, health, retirement, energy and the environment.

For much of the twentieth century, the working model of intertemporal choice was the (exponential) discounted utility model developed by Ramsey (1928) and Samuelson (1937), which features time-separable utility flows that are exponentially discounted: i.e., utility flows are discounted with the function δ^t , where δ is the discount factor and t is the horizon of the utility flow. This framework has many things to recommend it, including parsimony, generality/portability, and a single welfare criterion (which is implied by dynamically consistent preferences).¹ But the exponential discounting utility model is not descriptively accurate.

People seem to struggle when they make intertemporal tradeoffs, a phenomenon which has been extensively discussed by moral philosophers, political economists, psychologists, and policymakers. Writings about self-control and self-management are almost as old as written language itself.² As literature developed in the ancient world, the focus on self-control intensified. Greek philosophy contains many analyses about the challenges and virtues of self-management. Plato reports that Socrates described the soul as a charioteer (reason) with a pair of horses, one “noble” and the other one unruly and difficult to control.³ Aristotle emphasized the virtue of temperance and the human propensity to engage in self-defeating behaviors. “For moral excellence is concerned with pleasures and pains; it is on account of the pleasure that we do bad things and it is on account of the pain that we abstain from noble ones.” For Aristotle, an intemperate person has an appetite for pleasant things and chooses them at the cost of other, better things.⁴ Issues of self-control were explored by David Hume (“reason is the slave to the passions”) and Adam Smith (who distinguished between self-defeating “passions” and far-sighted “interests”).⁵ Much of the soft and hard paternalism reflected in the modern welfare state

¹ For another early formal analysis of intertemporal choice, see Koopsman (1960).

² For example, the ‘wisdom book’, *The Maxims of Ptahhotep* includes numerous recommendations for self-restraint. This text was likely written during the Old Kingdom or the Middle Kingdom of Ancient Egypt (Fox, 1983, dates the book to the 21 century BCE).

³ See *Phaedrus* sections 246a–254e. Also see discussion of in Chapter 7 of Nussbaum (2001).

⁴ See *Nicomachean Ethics, Book II Chapter 3 and Book III Chapters 10-11* (W. D. Ross translation).

⁵ On Hume, see Radcliffe (2018). On Smith, see Ashraf, Camerer, and Loewenstein (2005).

is an effort to influence intertemporal choices (e.g., see the discussion of Social Security in Feldstein 1985).

While much of economics in the mid-twentieth century modeled individuals as having a clear objective function and no self-control problems, research exploring self-control and intertemporal choice has blossomed in recent decades. Important milestones include two volumes published by the Russell Sage Foundation on intertemporal choice, edited respectively by Loewenstein and Elster (1992) and Loewenstein, Read, and Baumeister (2003). For a highly influential review of the development of the field see Frederick, Loewenstein and O'Donoghue (2002). Subsequent reviews include Chabris et al (2010) and Sprenger (2015). Cohen et al (2016) review the conceptual and methodological challenges associated with the measurement of intertemporal preferences. Cohen et al also discuss relevant neuroimaging evidence, which we completely omit from this current review.⁶

In this handbook chapter, we review the latest research on intertemporal choice and identify important open questions for our understanding of human behavior. We begin (Section 2) by examining formal models of intertemporal choice, because models provide a lens with which to examine empirical evidence and help identify new questions to explore. Most models we review share the unifying feature of giving some special priority to the present. To formalize the idea that the present is qualitatively treated differently than other periods, we introduce a meta-category of models: *present-focused preferences* exist if agents are more likely in the present to choose an action that generates immediate experienced utility, than they would be if all the consequences of the actions in their choice set were delayed by the same amount of time. More informally, this amounts to people choosing more impatiently for the present than they do for the future.

We intentionally use the term *present-focus*, rather than the more common term *present-bias*, because bias implies a prejudgment that the behavior is a mistake. Models that produce present-focused preferences include: hyperbolic and quasi-hyperbolic discounting (i.e., models with present bias); temptation that is experienced when choosing for now but not when choosing for the future; an interaction between myopic and planner selves; objective counter-party risks; and distortions in the perception of time or in forecasting the future. *Present-focused preferences* serve as a meta-category that identifies key commonalities among most of the models in the intertemporal choice literature. We also identify the key contrasts that differentiate the large number of present-focused models. Specifically, in Section 2.7 we provide a table that summarizes some of these differences.

In Section 3, we identify 10 key empirical regularities that have been well-documented in the literature: high required rates of return for money, higher required rates of return for consumption, preference reversals, procrastination, naiveté, large effect of transactions costs, a demand for commitment, the existence of paternalistic policies, and a preference for improving sequences. We accompany each of these 10 regularities with a closely related open question.

⁶ See Camerer et al (2015) for a related discussion of the literature on neuroeconomics.

1. Why do individuals have such a high required rate of return for money if impatience is fundamentally about consumption flows rather than financial flows?
2. How substitutable is consumption (and effort) across time? How does the substitution of consumption affect measured discount rates?
3. What are the key mechanisms that cause preference reversals?
4. Why do people underestimate their own procrastination? What are the relative roles of naiveté about present-focus versus other explanations for underestimating procrastination, such as overconfidence about the effort required and limited memory?
5. Why don't people learn and anticipate their present-focused behavior? What can help people correctly anticipate their future present-focused behavior?
6. Is the large effect of small transactions costs on behavior primarily related to present focus, or is it some other channel?
7. Why do households have such low levels of liquid net wealth, relatively high levels of illiquid net wealth, and high marginal propensities to consume out of liquid wealth changes?
8. Why is pure commitment so rare in markets and why is willingness to pay for commitment usually so low (even in lab experiments designed to measure the taste for commitment)?
9. What welfare criterion should we use to evaluate intertemporal tradeoffs?
10. How do we integrate models of discounting with the other factors leading to a preference for improving sequences? When does a preference for an improving sequence become important relative to discounting?

In section 4 we present six more open questions:

11. How soon is “now”? How fast does value decline over time?
12. What types of decisions involve temptation? How quantitatively important is temptation (as opposed to present-bias)?
13. How important are a variety of mechanisms for intertemporal choice, including probability weighting, trust, and heuristics, both in the lab and in the field? What role do alternative psychological conceptions play in intertemporal choice?
14. How stable are time preferences? How general are they across domains?
15. How malleable is time preference? How effective are self-management techniques?
16. Are households saving optimally for retirement?

Section five concludes.

2. Present-focused Preferences: Theoretical Commonalities

All animals, including humans, tend to pursue instant gratification, even when such immediate rewards are obtained by foregoing a substantially larger amount of delayed gratification. In section 3, we discuss both the qualitative and quantitative evidence (from the field and the lab) for this empirical regularity. In the current section, we describe several related theoretical frameworks that all generate a preference for immediate gratification. We focus on their similarities and accordingly group these conceptual frameworks together into a category that we call “present-focused preferences.”

Present-focused preferences exist if agents are more likely in the present to choose an action that generates immediate experienced utility, then they would be if all the consequences of the actions in their choice set were delayed by the same amount of time.⁷

Our definition compares an agent’s action in a situation in which they choose for the present to one in which they choose for the future (in a binding way)—for instance suppose eating ice cream delivered more immediate experienced utility than eating kale. Then, an agent displays present-focus if when choosing for today, they choose ice cream over kale, but when choosing for tomorrow, they choose kale over ice cream.

Circumstances that elicit such binding actions do not only arise in the laboratory or in artificial environments (though binding choices are easier to create in a controlled environment). Consider the person who tends to eat ice cream when it is immediately available, but tends to explicitly or implicitly choose not to eat ice cream in the future by not putting ice cream into his shopping cart at the supermarket.

Our definition of present-focused preferences refers to actions that generate immediate experienced utility—for instance, eating your favorite food, relaxing with friends, creating art or music, sexual activity, enjoyable hobbies, drinking alcohol, or recreational drug use.⁸ Increased immediate experienced utility can be derived not only from engaging in pleasurable or gratifying activities, but also by postponing activities that give immediate displeasure: consider a researcher who postpones working on a referee report whenever she can, but agrees to write the referee report in the first place (when the deadline is far away). In this example, the referee report generates instant displeasure (e.g., reading the paper is a slog), but generates long-term benefits (e.g., knowledge of the literature, or a reputation or self-image for being a good citizen).

Practical identification of activities that generate immediate experienced utility is conceptually challenging but can be approached in multiple ways. First, one could ask people which activities give them immediate pleasure, and even ask them to report that pleasure on a quantitative scale (see Wertenbroch 1998, Field Study 1, and related conceptualizations in Dellavigna and

⁷ By “more likely”, we mean weakly more likely and at least in some contexts strictly more likely.

⁸ See Loewenstein (1996) for a related discussion of visceral drives.

Malmendier 2004, 2006, and Oster and Morton 2005).⁹ Second, one could ask people which activities give *other* people immediate pleasure, both qualitatively and quantitatively. Third, one could measure willingness to pay for (marginal) experiences that only have immediate consequences: e.g., the willingness to pay for eating ice cream normally vs. having the same (unhealthy) ingredients safely ingested while one's taste buds are numbed. Fourth, one could measure neural activity during different activities and thereby identify activities that generate immediate neural activation in pleasure circuits (e.g., with neural probes in non-human animals – e.g., Schultz, Dayan, and Montague, 1997 – or with fMRI in humans – e.g., Rangel, Camerer, and Montague, 2008; Hare, Camerer, and Rangel, 2009).

The definition of present-focused preferences is intentionally unconstrained with respect to the nature/existence of self-control problems, issues that we will return to throughout this Section. Moreover, even though present-focused preferences generate dynamically inconsistent *choices*—that is, the choice between x and y made at date t for date t' is different from the choice between x and y at date t' -- they may result from dynamically consistent *preferences* (which we define below). *Preferences* can be dynamically *consistent* even if *choices* are dynamically *inconsistent*.

A standard definition of dynamic consistency in preferences is as follows: *Dynamic consistency in preferences between dates t and $t' > t$ arises when a person's state contingent preferences for actions taken at date t' , expressed at date t , are consistent with her state contingent preferences for actions taken at date t' , expressed at date t' . Dynamic inconsistency in preferences arises if there is any pair of values t and $t' > t$, which is not characterized by dynamic consistency in preferences.*¹⁰ In this definition “state contingency” incorporates all aspects of the choice. The state includes both internal phenomena like a twisted ankle (which would make exercise difficult and suboptimal) as well as external constraints, like the absence of a tempting good in a choice set. All of these contingencies are absorbed into the state contingency referred to in the definition of dynamic consistency of preferences.

The question of whether preferences are dynamically consistent can be applied to a model. One can ask what a particular model implies about the state-contingent preferences that are held at date t (for choices implemented at date t') and compare them to the state-contingent preferences that will be held at date t' . Likewise, the definition can be practically deployed as a series of hypothetical questions asked at date t : “What would you like to choose at date t' conditional on state X , knowing that you are not binding yourself to this answer?”

As we will emphasize below, some theories that generate present-focused preferences feature dynamically inconsistent preferences (e.g., present bias, which is discussed in Section 2.1). However, other theories that generate present-focused preferences feature dynamically consistent preferences (e.g., temptation models, which are discussed in Section 2.2).

⁹ See the large literature on subjective well-being, including experiments that triangulate choice data with self-reported data on tastes (e.g., Hare, Camerer, and Rangel 2009; Kahneman and Deaton 2010; and Stevenson and Wolfers 2013). These measures have also been used in practical domains. For example, most hospitals now have pain scales posted on the walls, so that patients can identify the level of discomfort that they are experiencing.

¹⁰ Some types of dynamic inconsistency arise in domains that are unrelated to time preferences (e.g., Andreoni, Aydin, Barton, Bernheim, and Naecker 2016).

We will also discuss the issue of commitment. Some theories that generate present-focused preferences feature a taste for commitment: *a strictly preferred restriction of one's own future choice set*. Here we are interested only in the intra-personal taste for commitment – e.g., a smoker who is trying to quit and chooses to flush her cigarettes down the toilet. We exclude any inter-personal strategic reasons for commitment – e.g., a general who burns a bridge *behind* her army, signaling to the opposing forces that her soldiers will fight to the death now that retreat is not an option. In this chapter, when we discuss the phenomenon of commitment we refer exclusively to the taste for (or choice of) commitment that arises from pure, intra-personal mechanisms.

We conclude the chapter by presenting a 2×2 table, which reviews all of the theories of present-focused preferences and organizes them along two dimensions: dynamic consistency of preferences and a taste for commitment. We will end up filling all four boxes with theories that fall under the broad rubric of present-focused preferences. We first discuss these theories in isolation, before drawing them together at the end of section two.

2.1. Present-biased preferences

The name *present-focused* preferences is a variation (and generalization) of the concept of *present-biased* preferences. Present-focused preferences include both present-biased preferences and many other related models. Present-biased preferences are the most commonly used intertemporal choice model in behavioral economics, so we describe it first.

Present bias (Laibson 1997, O'Donoghue and Rabin 1999) is also referred to as *quasi-hyperbolic discounting*, which highlights its intellectual debt to the earlier literature on hyperbolic discounting (e.g., Loewenstein and Prelec 1992).¹¹ Present-biased preferences are expressed:

$$U_t = u_t + \beta\delta u_{t+1} + \beta\delta^2 u_{t+2} + \beta\delta^3 u_{t+3} + \dots \quad (1)$$

Here U_t is total utility, u_t is flow utility in period t , β is the present bias parameter, and δ is the long-run discount factor. Phelps and Pollak (1968) first used this framework to model *intergenerational* preferences, in contrast to the intra-personal preferences that we discuss now. For Phelps and Pollak, u_t is the utility of generation t and u_{t+1} is the utility of generation $t+1$. In the behavioral economics literature, the discounting is intra-personal instead of being inter-generational. Accordingly, u_t is the utility flow that an agent experiences during period t and u_{t+1} is the utility flow that the same agent experiences during period $t+1$.

In this framework (like most other intertemporal choice models in economics) the object being discounted is a stream of utility flows, which are distinct, in principle, from financial flows. In other words, the theory is about how agents discount pleasures and pains experienced at particular points in time, and not about how agents think about the timing of financial events

¹¹ For earlier examples from the psychology literature, see Herrnstein (1961) and Ainslie (1974). For a different, but conceptually related functional form, see Benhabib, Bisin, and Schotter (2010). For a closely related formulation based on salience, see Akerlof (1991).

(e.g., whether my fully anticipated paycheck is auto-deposited in my bank account on Friday or Monday -- assuming that the interest rate is zero and my short-term liquidity is more than adequate to cover all of my expenses and other transactions over the weekend). The distinction between utility discounting and money discounting is reviewed in Sections 3.1 and 3.2, as well as in Cohen et al (2016). To our knowledge this distinction was first clearly articulated by Ramsey (1928): “This rate of discounting future utilities must, of course, be distinguished from the rate of discounting future sums of money. If I can borrow or lend at a rate r I must necessarily be equally pleased with an extra £1 now and an extra £(1 + r) in a year's time, since I could always exchange the one for the other. My marginal rate of discount for money is, therefore, necessarily r , but my rate of discount for utility may be quite different, since the marginal utility of money to me may be varying by my increasing or decreasing my expenditure as time goes on.”

To illustrate the importance of the timing of utility flows, consider the example of exercise (e.g., gym attendance). In most applications, exercise is assumed to have immediate logistical, and psychic costs and delayed health benefits (e.g., DellaVigna and Malmendier 2004, 2006), which will engender a self-defeating pattern of low levels of exercise, even among those who pay for gym memberships. However, the opposite self-regulation problem might apply to people who greatly enjoy the physical sensation of exercising. These presumably rare types will tend to exercise too much, for the same reasons (in the model) that the rest of us watch too much television and eat too many donuts. Modelers need to be careful in the assumptions that they make about the timing of utility flows (for example, see the related analysis in Augenblick, Niederle and Sprenger, 2015) and the valence of utility flows.

In this model it is natural to place bounds on the discounting parameters: $0 < \beta \leq 1$ and $0 < \delta \leq 1$. In the literature, the notation distinguishes the term ‘discount factor’, which refers to a specific weighting parameter β or δ , and ‘discount function’, which refers to the function of horizon-dependent utility weights, which, for present-biased preferences, is given by

$$D(t) = \begin{cases} 1 & \text{if } t = 0 \\ \beta\delta^t & \text{if } t \geq 1 \end{cases} . \quad (2)$$

In Equations (1) and (2), β is not exponentiated (whereas δ is). Because of this difference, it is helpful to rewrite these preferences as

$$U_t = u_t + \beta[\delta u_{t+1} + \delta^2 u_{t+2} + \delta^3 u_{t+3} + \dots] . \quad (3)$$

Written this way, it is clear that present biased preferences embed a continuation payoff stream that is weighted by β and then exponentially discounted thereafter. When $\beta = 1$ these preferences revert to exponential discounting. (As a general point, it is useful when behavioral models nest the classical model as a special case.) In typical calibration/estimation of these discounting factors, $\beta \ll 1$ and $\delta \cong 1$.

Present-biased preferences are *present-focused*. Under present-biased preferences, $\beta < 1$ engenders a preference to experience pleasurable activities in the present instead of the future. For example, assume that exercise has an immediate cost of c and a delayed (health) benefit of b .

Assume that all of these utility flows are differences relative to an alternative activity which we will call napping. To further simplify this example, assume that $\delta = 1$. If $b > c > \beta b$, then the present-biased agent prefers to nap today and to commit to exercise tomorrow:

$$-c + \beta b < 0,$$

$$\beta(-c + b) > 0.$$

Present-biased preferences are dynamically inconsistent: e.g., in this example, the agent at time t prefers to exercise at time $t' > t$ and the agent at time t' prefers to nap at time t' .

Present-biased preferences engender a (potential) preference for commitment if the agent has sophisticated or partially sophisticated beliefs (see Strotz 1955, Phelps and Pollak 1968, Laibson 1997, and O'Donoghue and Rabin 1999; however, see Laibson 2015 for reasons that commitment might still not arise in practice even for partially or fully sophisticated agents). Sophisticated beliefs imply that the agent has a correct theory of her own future time preferences.

Partially sophisticated beliefs (i.e., partially naïve beliefs) imply that she recognizes that she is present-biased, but she believes that her future selves ($t' > t$) will each have a higher value of β than those future selves will actually have $\beta(t') < E_t[\beta(t')] < 1$ (O'Donoghue and Rabin 1999). In the literature, $\hat{\beta}$ is used as a short-hand for $E_t[\beta(t')]$. Fully naïve beliefs imply that the agent believes that her future selves will not be present biased at all: $\hat{\beta} = 1$. Accordingly, fully naïve beliefs imply that a present-biased agent won't have a preference for commitment. She believes that all future selves will act according to her current preferences for their future actions.

Present-biased preferences are a special case within a large class of discounting models that are distinguished by high discount rates in the short-run and low discount-rates in the long-run – i.e., a monotonically falling discount rate (see Strotz 1955). The discount rate is the local rate of decline in the discount function, $D(t)$. In discrete time models, the local rate of decline of the discount function (per unit of time) is

$$-\frac{D(t+1) - D(t)}{D(t)}. \quad (4)$$

In continuous time models (with differentiable discount functions), the local rate of decline of the discount function is

$$-\frac{D'(t)}{D(t)}. \quad (5)$$

Under exponential discounting, the discrete-time discount rate is $1 - \delta$ and the continuous-time discount rate is $-\ln \delta \approx 1 - \delta$.¹² These discount rates are constant, and therefore not horizon dependent.

¹² This approximation represents a first-order Taylor expansion.

Contrast this with the case of present bias. Under present-biased discounting in discrete time, the short-run discount rate is $1 - \beta\delta$ and the long-run discount rate is $1 - \delta < 1 - \beta\delta$.

There are also generalizations of present-bias in continuous time, which we omit in this short survey (see Harris and Laibson (2012)). Continuous-time implementations generate desirable modelling properties, like smooth policy functions, equilibrium uniqueness, and—like the exponential discounting model—e a single value function that the individual uses to judge their own welfare.

2.2. Unitary-self models with temptation

Present-biased preferences induce dynamically inconsistent preferences. As such preferences began to gain influence in the 1990s, some economists pointed out that many of the phenomena that were being explained by models with dynamically inconsistent preferences, could also be explained by models with rational beliefs and dynamically consistent preferences (e.g., Dekel, Lipman and Rustichini, 2001; Gul and Pesendorfer, 2001; Laibson 2001; Bernheim and Rangel, 2004; Dekel, Lipman and Rustichini, 2009; Gul and Pesendorfer, 2004; Noor, 2007; Noor, 2011; Lipman and Pesendorfer, 2013). Some of these unitary-self models are called *temptation models*. (A related psychology literature also studies temptation effects: e.g., Baumeister et al, 1998; Muraven, Tice, and Baumeister, 1998; Muraven and Baumeister, 2000.) Present-biased models also feature temptation properties, though the nature of the temptation is different, as we will explain below.

Unitary-self models of self-control problems come in many different particular forms. They are related by the principle that choice *sets* affect utility, including the *options that are not chosen* (see Kreps, 1979). Consider an example in which a decision-maker at time t can, in principle, engage in two mutually exclusive activities during a future period $t' > t$: exercise or nap. Assume that exercise has immediate consequences that are hedonically aversive, and therefore exercise is not tempting. By contrast, assume that nap has immediate consequences that are pleasurable, and therefore nap is tempting. Suppose that the decision-maker must choose her choice set for period t' during t . At period t , the decision-maker has preferences over three possible choice sets for period t' : $\{\text{exercise, nap}\}$, $\{\text{exercise}\}$, $\{\text{nap}\}$. A singleton choice set implies that the decision-maker must choose that action (there is no opt-out).

Two scenarios illustrate some of the key ideas in this literature: “commitment with strong temptation” and “commitment with weak temptation.” For all of the analysis in this subsection, we will assume that there is no risk/uncertainty, thereby *removing* the classical option-value reason for preferring flexibility (however, see Kreps 1979, for an alternative framework in which agents have a preference for flexibility for its own sake).

Commitment with strong temptation: The period t preferences over choice sets available at time t' , could take the following form:

$$\{\underline{\text{exercise}}\} > \{\text{exercise, } \underline{\text{nap}}\} \sim \{\underline{\text{nap}}\}. \quad (6)$$

To simplify exposition, we underline the outcome that occurs at period t' conditional on each choice set that could be chosen in period t . In this scenario, during period t the decision-maker prefers that at period t' she will exercise without the temptation to nap. Therefore, she chooses to tie her own hands by compelling herself to exercise at time t' . If the choice set {exercise, nap} were available at period t' , then nap would be chosen (because nap is the tempting option and, in the current scenario, this temptation is *strong*). Consequently, at time period 0, the agent is indifferent between the (future) choice sets {exercise, nap}, which will induce the tempting action of napping, and the restrictive choice set {nap}. In either case, she'll nap. (This is the same type of temptation that emerges in the model of present bias.)

Commitment with weak temptation: The period t preferences over choice sets available at time t' , could take the following form:

$$\{\underline{\text{exercise}}\} > \{\underline{\text{exercise}}, \text{nap}\} > \{\underline{\text{nap}}\}. \quad (7)$$

At date t , the decision-maker prefers that at time t' she will exercise without the temptation to nap (as before). Therefore, she chooses to tie her own hands by compelling herself to exercise at time t' . If the choice set {exercise, nap} were available at time period t' , then exercise would still be chosen. Even though nap is the tempting option, this temptation is weak, so exercise would still be chosen at period t' . But the agent does pay some utility cost when she rejects a tempting option *that is in her immediate choice set*, so at date t , equation (7) features two strict inequalities, reflecting the costs of temptation.

In these analyses, temptation plays two roles. First, it can affect the final choice, as it does in the first example (strong temptation). The availability of the option to nap causes nap to be chosen in period t' (even though it is preferred to remove that option during period t so it isn't in the consideration set at period t').

Second, temptation can affect welfare, even when it does not affect the final choice. In the second example (weak temptation), exercise is chosen whenever it is available. However, the decision-maker still strictly prefers to remove nap from her choice set. In this scenario, nap is a temptation that makes the decision-maker worse off, even though she never chooses it when exercise is available. Intuitively, it is psychologically costly to resist a temptation.¹³

Let's now generalize this example. Consider an agent at time t who has two *immediate* choices to make from *unrestricted* choice sets: do I choose to exercise or nap at time t and do I choose (now at time t) to exercise or nap at time t' . The choices made at time t for time t' will be implemented by choosing singleton choice sets for time t' , thereby committing actions at time t' .

¹³ Though rarely discussed, temptation effects could, in principle, have effects opposite in sign to those discussed in this subsection. If a decision-maker takes pride/pleasure in successfully resisting temptations, then a third case could arise. *Self-satisfaction when resisting temptation:* The period 0 preferences over choice sets available at time 1, could take the following form: $\{\underline{\text{exercise}}, \text{nap}\} > \{\underline{\text{exercise}}\} > \{\underline{\text{nap}}\}$. At date 0, the decision-maker prefers that at time 1 she will exercise and she foresees that she will take extra pleasure in exercising if she does so by using her own willpower, instead of relying on an extrinsic commitment mechanism. See Kreps (1979).

Under strong temptation, the agent will choose to nap at time t and will commit to exercise at time t' , implying that these preferences are present-focused.

Under weak temptation, the agent will choose to exercise at time t and will also commit to exercise at time t' . Moreover, at time t , the agent will *strictly* prefer (i.e., she is willing to pay) to restrict her choice set at time t' , even though this commitment does not affect her choice at time t' . This phenomenon – willingness to pay for commitments that do not actually affect one's choices – is not predicted by present-biased preferences (see 4.2 for evidence of such commitments). Present-biased agents may be willing to pay for commitment, but only if that commitment changes their choices.

Summing up, the unitary-self model with temptation features dynamically consistent preferences (all selves have aligned state-contingent preferences, where the contingency includes the relevant choice set). The unitary-self model also features commitment, unless the agent is naïve about her preferences (e.g., see Ahn, Iijima, and Sarver 2017).

2.3. Multiple-self models with simultaneous selves.

A different branch of present-focused models is based on the idea that competing sets of interests simultaneously pull the decision-maker in different directions (e.g, Thaler and Shefrin 1981; Shefrin and Thaler, 1988; Hoch and Loewenstein 1991; Loewenstein, 1996; Bernheim and Rangel, 2004; Loewenstein and O'Donoghue, 2004; McClure et al 2004, 2007; Fudenberg and Levine, 2006, 2011, 2012; Brocas and Carrillo, 2008a, 2008b, 2012; Jackson and Yariv, 2014, 2015).

This conceptualization of simultaneously competing internal interests can be traced back to Greek philosophers and probably has even earlier precedents. In the field of economics, the idea originates with Smith (1759),¹⁴ who adopted a framework that echoed classical sources with which he would have been familiar (e.g., Plato's discussion of Socrates, referenced in the introduction). Smith contrasted two sets of motivational systems: passions, which tend to be myopic, and a far-sighted internal spectator who attempts to rein in the free expression of our passions. "The spectator does not feel the solicitations of our present appetites. To him the pleasure which we are to enjoy a week hence, or a year hence, is just as interesting as that which we are to enjoy this moment." Related models reappeared in the economics literature with the contributions of Thaler and Shefrin (1981) and Shefrin and Thaler (1988). They hypothesized competition between a myopic "doer" and a far-sighted (and dynamically consistent) "planner."

Subsequent contributions (e.g., Fudenberg and Levine 2006, 2011, 2012) have used formal frameworks that populate a middle-ground between the planner-doer model of Thaler and Shefrin and the unitary-self models described in the previous subsection.¹⁵

¹⁴ See Ashraf, Camerer, and Loewenstein (2005) for a fascinating review of the rich veins of behavioral economics that can be found in the work of Adam Smith, especially *The Theory of Moral Sentiments* (1759).

¹⁵ "While we find the language of multiple 'selves' to be suggestive, the model can equally well be interpreted as describing the behavior of a single 'self' whose overall behavior is determined by the interaction of two subsystems" (Fudenberg and Levine, 2006, p. 1450).

Fudenberg and Levine begin their 2006 paper with a quote from McIntosh (1969): “The idea of self-control is paradoxical unless it is assumed that the psyche contains more than one energy system, and that these energy systems have some degree of independence from each other.” Fudenberg and Levine model two types of selves that play a repeated stage game: a patient (dynamically consistent) long-run self and a sequence of myopic short-run selves. These two types of selves share the same preferences over the immediate outcomes of each current stage-game. However, only the long-run self cares about payoffs in future stage games. Fudenberg and Levine assume that each stage game is played in two phases. In the first sub-period of the phase game, the long-run self chooses to exert a level of (costly) self-control that influences the preferences of the current short-run self. In the second sub-period of the stage game, the short-run self takes the final decision.

Models with simultaneously competing selves also fit into the general class of present-focused models. Here too, (unconstrained) agents are relatively more likely to choose hedonically pleasurable activities for the present, then they are to commit to choose the same hedonically pleasurable activities for the future. For example, in the model of Fudenberg and Levine, the current short-run self wants to experience immediate pleasures, making it difficult (i.e., costly) for the long-run self to dictate the outcome. In equilibrium, unconstrained choices for the present will be biased toward immediate gratification because the long-run self won’t be willing to exert enough costly self-control to force the hand of the myopic short-run self. However, when the same person makes binding commitments for the future, the current short-run self doesn’t care about what will happen, so it can be easily manipulated by the long-run self. Accordingly, choices made for the future will be less prone to favor short-run pleasures.

Multiple-self models with simultaneous selves feature a preference for commitment. The long-run self typically wants to restrict future choice sets to make it easier to resist the drives of future short-run selves.

It is unclear how to categorize multiple-self models with simultaneous selves with respect to the property of dynamic consistency of preferences. The long-run selves in these models are (usually) dynamically consistent in their preferences. The long-run selves may *appear* to be dynamically inconsistent because of their fraught interactions with their myopic selves. The myopic short-run selves create drives that the long-run selves need to resist or, if that is very costly, accept. Once one properly accounts for these myopic psychological drives, and the psychic costs that must be paid to overcome them, the long-run selves have dynamically consistent preferences (just like the unitary self in the temptation models).

Nevertheless, it would be odd to say that the *full* set of preferences is dynamically consistent in multiple-self models with simultaneous selves. These models feature competing preferences (e.g., the preferences of long-run and short-run selves), which are not aligned.

2.4. Objective risks that reduce future value

Objective risks can create another source of present-focused preferences. Assume that future rewards may be ‘lost’ before they are received – e.g., the counter-party making the promise of the future payment could go bankrupt or skip town. If the hazard rate of such losses is

horizon dependent, $\rho(\tau)$, then a perfectly patient decision-maker should discount rewards at rate $\rho(\tau)$ at horizon τ . If this hazard is constant, discounting will appear exponential, but a changing hazard can generate discount functions with a hyperbolic shape (see Sozou 1998; Azfar 1999; Weitzman 2001; Halevy 2005; Dasgupta and Maskin 2005; Fernández-Villaverde and Mukherji 2006; Halevy 2014, 2015).

Hyperbolic discounting can be generated by uncertainty about what the hazard rate is (even if it is a constant rate). For instance, Sozou (1998), Azfar (1999), and Weitzman (2001) assume that $\rho(\tau)$ is not horizon-dependent, so $\rho(\tau)$ is a constant, and ρ is not known because it has been drawn from a distribution. Consequently, at every point in time the true (realized) value of ρ has a posterior distribution that is changing. The more time that passes without a loss, the more likely that one of the low values for ρ was originally drawn. This generates intertemporal choices that are equivalent to those that would be generated by a discount rate that declines as the horizon increases (e.g., hyperbolic discounting).¹⁶

This is another example of present-focused preferences arising from an agent whose preferences are dynamically consistent. In this scenario, the agent will look as if she has hyperbolic time preferences (in the sense that her discount rate is declining with the horizon), but she won't have a self-control problem, she won't have dynamically inconsistent preferences, and she won't view commitment as something valuable. She will always prefer more choice to less (all else equal).

This is an illustration of a more general point: exponential discounting is sufficient, but not necessary for dynamically consistent preferences. If all selves agree on the same (non-stationary) non-exponential discount function, then they will have dynamically consistent preferences.

2.5. Models with psychometric distortions

Models with psychometric distortions posit that time and risk are perceived with psychological distortions that tend to generate present-focused preferences. For example, suppose time is perceived with a concave (subjective) transformation (as argued by Read 2001, Takahashi 2005, Ebert and Prelec 2007, and Zauberman et al. 2009). To fix ideas, suppose that agents discount with an exponential discount function, $D(\tau(t)) = \delta^\tau$, but that they perceive objective time, t , with a concave transformation $\tau(t)$. Then the discount rate inferred by the social scientist (with respect to objective time), will be given by

$$-\frac{D'(\tau(t))}{D(\tau(t))} \tau'(t) = (-\ln \delta) \tau'(t).$$

Note that when $\tau'(t) = 1$, so that subjective and objective slopes align, the researcher will simply impute a constant discount rate, $-\ln \delta$. However, if τ is a concave function, the imputed

¹⁶ Dasgupta and Maskin (2005) show that while Sozou's model cannot explain preference reversals, uncertainty about when payoffs will be realized can.

discount rate is declining as the objective horizon, t , increases. To see this, take the derivative of the imputed discount rate with respect to the discounting horizon:

$$\frac{d[(-\ln \delta)\tau'(t)]}{dt} = (-\ln \delta)\tau''(t) < 0.$$

Hence, the agent's psychometric distortion in the perception of time produces behavior that mimics the behavior that one observes when the discount function has a declining discount rate in the (objective) horizon of the discounted reward. Relatively high short-run (effective) discount rates and relatively low long-run (effective) discount rates generate present-focused preferences for the same reasons that we see them in section 2.1.

This mechanism is one microfoundation for hyperbolic discounting. However, this version of hyperbolic discounting is different from the one implied by present bias. For example, this psychometric distortion does not seem to imply a motive for commitment. Rather a psychometrically distorted view of the future is like an optical illusion. When you become aware of the optical illusion, you try to correct your misperception (i.e., make some effort at debiasing), rather than restricting your future choice sets. On the other hand, psychometric distortions do tend to generate dynamic inconsistency in preferences: an agent with distorted time perception has preferences for future tradeoffs which are inconsistent with the preferences that she will actually have in the future.

In the model above, agents perceive objective time with a distortion, leading to present-focus. Agents might instead perceive *delays* between two options with distortions, such that the delay between t' and t is perceived as $\tau(t' - t)$ rather than $\tau(t') - \tau(t)$. This psychometric distortion of delays can be one source of *subadditivity* (Read 2001; Scholten and Read 2006; see also Glimcher, Kable, and Louie 2007; and Kable and Glimcher 2010).¹⁷ Specifically, subadditivity arises when discounting over a time interval is greater when the interval is divided into subintervals, which is a robust empirical regularity. Subadditivity produces choices that mimic a key property of hyperbolic discounting: diminishing discount rates as the length of the discounting period is increased. Despite this similarity with hyperbolic discounting, discounters who exhibit subadditivity will not choose commitment and their preferences will be dynamically consistent.¹⁸ Neither perceiving delays with distortions nor subadditivity generate present-focused preferences, since an additional delay is perceived the same regardless of whether it begins now or in the future. As a result, a sub-additive agent will make the same choice today (e.g., 1 util now vs. x utils tomorrow) as she will for the future (1 util in t days vs. x utils in $t + 1$ days).¹⁹

¹⁷ Subadditivity may be rational if there is implicit information (e.g., about riskiness) in the way that information is elicited from subjects. For example, it may be that any delay in payment is associated with an implied risk, and the magnitude of the risk is not (highly) dependent on the length of the delay. Under this interpretation, subadditivity could be included in subsection 2.5.

¹⁸ Specifically, their views about tradeoffs between utils at t and utils at $t' > t$ does not depend on when (on or before date t) they are asked to express this preference.

¹⁹ Cohen, Ericson, Laibson, and White (2016) discuss the evidence for two distinct features of hyperbolic discounting: discount rates fall as a front-end delay is added and as the length of the discounting period is increased. Subadditivity can account for the latter regularity.

Other psychometric distortions produce present-focused preferences. For example, subjective probability distortions, such as the probability weighting function of Kahneman and Tversky (1979), will produce behavior that mimics hyperbolic discounting. Specifically, if the present is perceived to be risk free (in the sense that promised rewards will be delivered with certainty), but the near future is perceived to introduce a tiny amount of risk, then the existence of a certainty effect (whereby a very small amount of risk is subjectively treated like a large reduction in likelihood), will generate behavior like present bias. For related arguments see the non-expected utility frameworks in Prelec and Loewenstein (1991), Quiggin and Horowitz (1995), Keren and Roelofsma (1995), Weber and Chapman (2005), Halevy (2008), Epper, Fehr-Duda, and Bruhin (2011), Baucells and Heukamp (2012), Andreoni and Sprenger (2012), Epper and Fehr-Duda (2015), Chakraborty (2017). These models do not generate a demand for commitment, but they do generate dynamically inconsistent preferences.

Heuristic reasoning can also generate present-focused preferences. Proportional reasoning can lead to behavior that has many of the properties generated by a hyperbolic discount function. For example, if a delay from 0 days to 1 day seems like a large delay (because the average delay is $\frac{1}{2}$) but a delay from 100 days to 101 days seems like a small delay because the average is $100\frac{1}{2}$, then agents may respond more aversively to rewards that are delayed in the former than in the latter case (Rubinstein, 2003; Read et al 2013; Ericson et al 2015). Agents that are using heuristic reasoning will not choose commitment, but they will express dynamically inconsistent preferences.

The focusing model of Kőszegi and Szeidl (2013) also generates present focused preferences in some circumstances. In this model, consumers choose among multi-dimensional consumption vectors, where each dimension represents an attribute. The consumer maximizes focus-weighted utility, which disproportionately weights attributes in which her options generate a greater range of consumption utility. Their model produces no present-focus when simply trading off utility at two different dates. However, agents may display present-focus when they frame a choice (like exercising today) as an isolated current decision with a salient up-front cost (going to the gym) and a stream of tiny non-salient benefits (being slightly healthier, each day over the next few years because you exercised once). When instead framing the decision as a continuous stream of exercise (in the future) and a stream of *large* health benefits (that arise from a consistent routine of exercising), the agent makes a more patient choice. The all-in or all-out decision leads to patient choices, because now the streams of costs and benefits are both large and the focusing bias does not make the cost of exercise disproportionately salient. Agents with focusing bias will not choose commitment if the bias is a perceptual error, but they will express dynamically inconsistent preferences.

2.6. Models of myopia

The concept of myopic decision-making has been discussed for millennia. Though it is not possible to locate the intellectual origin of this idea there are many historical milestones.²⁰ For

²⁰ See Loewenstein (1992) for a review of the 19th and 20th century intellectual history of theories of intertemporal choice.

example, Böhm-Bawerk (1890) wrote that “we possess inadequate power to imagine and to abstract, or that we are not willing to put forth the necessary effort, but in any event we limn a more or less incomplete picture of our future wants and especially of the remotely distant ones. And then, there are all of those wants that never come to mind at all.”

Pigou (1920) similarly observed “that our telescopic faculty is defective, and that we, therefore, see future pleasures, as it were, on a diminished scale. That this is the right explanation is proved by the fact that exactly the same diminution is experienced when, apart from our tendency to forget ungratifying incidents, we contemplate the past.”

Even if we stipulate that people are “myopic,” it is not at all clear what that formally means. Loosely speaking, myopia is a failure to clearly see the future when it is forecastable in principle (e.g., if one “took their head out of the sand” or “stopped living completely in the moment”).²¹

Gabaix and Laibson (2017) model myopia as the consequence of *cognitive* noise that reduces the accuracy of signals about future events (see also related work by Commons et al 1982, 1991). Because agents know that their signals are noisy, they shade those signals toward their prior, causing their forecasts to be relatively unresponsive to true variation in future rewards. Consequently, future utility flows are given less weight than they would have if perceptions were not noisy. This implies that the agent will exhibit present-focused preferences: the agent’s choices will reflect greater sensitivity to the attributes – positive and negative – that are experienced in the present.

The model of Gabaix and Laibson produces behavior that mimics much of the behavior that arises under hyperbolic discounting. However, the model generates no taste for commitment because the agents are not encumbered by a self-control problem; their only problem is that they are unable to see clearly into the future. The model assumes that agents have *underlying* dynamically consistent preferences. However, the agent may not be able to act on these preferences because her noisy perceptions of future payoffs undermines the expression of these preferences, leading her to *exhibit* preferences that appear to be dynamically inconsistent to an outside observer who doesn’t appreciate her perceptual limitations.

Other forms of myopia also produce behavior that appears to be characteristic of dynamically inconsistent preferences. For example, one could interpret the focusing model of Kőszegi and Szeidl (2013) as a model of selective myopia (i.e., overlooking the attributes that are not focal). This model implies no commitment and generates (expressed) preferences that will appear to be dynamically consistent to an outside observer. See Frederick (2005) for another mechanism that can produce myopia: a low propensity to engage in cognitive reflection. See also Steele and Joseph (1990) for alcohol-induced myopia.

2.7. Overview of models of present-focused preferences

We have reviewed a large family of models that feature present-focused preferences. These models share the property that agents choose more impatiently for the present than when they

²¹ Early formulations of myopia were developed by Brown and Lewis (1981) and Jéhiel (1995).

choose for the future. Commonalities of models in the intertemporal choice literature have been noted by many authors (e.g., Barro, 1999; Krusell, Kuruşçu, and Smith, 2010, Gustman and Steinmeier 2012, Gabaix and Laibson 2017).

Despite this core similarity, the models in this literature have many contrasting properties, some of which are summarized in Table 1 below. The taste for commitment (which is captured in the rows of the table) is a particularly important differentiator, because this is directly observable in choices. In other words, we can (in principle) observe this smoking gun if we can find settings in which commitment is purely choice-set restricting and not confounded by other benefits (like tax advantages in illiquid retirement savings plans). However, see Laibson (2015, 2018) for reasons that we may not expect to see commitment in practice. For example, exogenous uncertainty may make commitment undesirable even for fully sophisticated present-biased agents who would commit themselves in an idealized world with no uncertainty.

Dynamic consistency of preferences is captured in the columns of Table 1. Dynamic consistency in preferences is not measurable by observing (binding) choices. Rather it is a property of a model. As an empirical object, it is measured by asking people about their current state-contingent preferences (in future states).

Table 1: Present-focused models categorized by two properties: commitment (rows) and dynamic consistency of preferences (columns).

	Dynamically Consistent Preferences	Dynamically Inconsistent Preferences
Commitment	<ul style="list-style-type: none"> •Unitary-self temptation models: 2.2 •Long-term self in multiple-self models: 2.3 	<ul style="list-style-type: none"> •Present-bias with partial sophistication: 2.1 •Other forms of hyperbolic discounting: 2.1
No Commitment	<ul style="list-style-type: none"> •Exponential discounting: 2.1 •Objective risks (non-exponential discounting): 2.4 •Myopia: 2.6 	<ul style="list-style-type: none"> •Present-bias with perfect naiveté: 2.1 •Psychometric distortions: 2.5 •Myopia: 2.6

Note: Each model is discussed in the associated subsection.

2.8. Models that do not generate present-focused preferences

Finally, we conclude Section 2 by emphasizing that there are other influential models of intertemporal choice that do not fall into the category of present-focused preferences. Present-focused preferences undergird a large part of the intertemporal choice literature, but not all of it.

Habits and related reference-point effects lead agents to experience larger utility flows when they can favorably compare current consumption to their own past consumption, to consumption of their peers, or to consumption that they had expected (e.g., Ryder and Heal, 1973; Becker and Murphy 1988; Loewenstein 1988; Abel, 1990; Hoch and Loewenstein, 1991; Loewenstein and Sicherman, 1991; Loewenstein and Prelec, 1992, 1993; Prelec and Loewenstein, 1998; Campbell and Cochrane, 1999; Laibson, 2001; Koszegi and Rabin, 2009). Some of these models generate

present-focused preferences, as in the preferred calibration of Koszegi and Rabin (2009). But most of these theoretical mechanisms have the effect of leading agents to choose consumption profiles with a relatively higher slope (i.e., less consumption now and more consumption in the future), so the comparisons to past consumption will be more favorable over the lifecycle.

Anticipation effects also encourage consumers to move rewarding events into the future so they can be savored in advance (Loewenstein 1987; Caplin and Leahy 2001, Kreps 1998). This idea has a long history. Loewenstein (1987) quotes Alfred Marshall on the topic: “When calculating the rate at which a future benefit is discounted, we must be careful to make allowance for the pleasures of expectation.” (Marshall, 1891, p. 178).

Finally, there is a class of models that have implications for the timing of the resolution of uncertainty (e.g., Kreps and Porteus 1978; Epstein and Zin 1989; Weil 1990). Depending on the calibration of these models, they can either rationalize a preference for early or late resolution of uncertainty.

3. Empirical Regularities and Open Puzzles

We turn now to a series of well-established empirical regularities, each of which is also associated with some important open questions. Additionally, we identify a number of additional open questions for future research.

3.1. Preferences over Monetary Receipt Timing

Empirical Regularity #1: Individuals require a high rate of return (RRR) for money earlier versus later tradeoffs.

One of the primary ways time preference has been measured has been with Money Earlier or Later (MEL) experiments, in which subjects choose between X dollars at an early date or Y dollars at a later date.²² While the models we discuss above are—almost always—fundamentally about the discounting of *utility*, it is hard to directly give an individual a unit of utility. Money is easily measured, is socially important, and is a placeholder for reward (albeit one that can be easily intertemporally shifted). Accordingly, a large literature examines how individuals choose between money earlier and later.

To describe the indifference points generated by individuals making choices over financial flows, we follow Cohen, Ericson, Laibson, and White (2016) and study the required rate of return (RRR). If an individual is indifferent between $\$x_1$ at time t_1 and $\$x_2$ at a later time t_2 , we define the RRR between time t_1 and t_2 to be $\ln \frac{x_2}{x_1}$. (This is approximately the percentage difference between the rewards received in the respective periods.) Likewise, we define the annualized RRR as²³

²² This section and Section 2 draw in part on Cohen, Ericson, Laibson, and White (2016).

²³ These definitions are derived from the following implicit definition of an instantaneous annual rate of return, r :
 $\exp(r[t_2 - t_1]) = \frac{x_2}{x_1}$.

$$\frac{1}{(t_2 - t_1)} \ln \frac{x_2}{x_1},$$

where time units are measured in years. For comparability, it will often be convenient to annualize the RRR, though this hides some important information about the type of choices from which the RRR was estimated.

Note that the RRR is not a preference parameter. The discount rate for *utility* requires additional assumptions on when and how financial flows translate into utility: does the receipt of money correspond to receipt of utility at that time? Does utility increase approximately linearly in the amount of money received, or is there substantial curvature even over small amounts of money? Cohen, Ericson, Laibson, and White highlight the difficulty of interpreting what MEL experiments imply for discounting of utility, and present different models that are used to translate MEL choices into (estimated) discount functions.

The first published MEL paper is Thaler (1981), who showed a number of anomalies using hypothetical MEL questions, including the key finding that the annualized RRR declines as the time horizon gets longer (holding the earlier payment date fixed).²⁴ Kirby (1997) used incentive compatible (sealed second bid auctions) to measure the RRR at different time horizons, and found similar results, finding that a hyperbolic function fit the data better than the exponential discount function. The MEL design was used by several other economists after Thaler (e.g. Loewenstein & Thaler, 1989; Prelec & Loewenstein, 1991); it was popularized in the psychology literature by Kirby and Herrnstein (1995). MEL experiments now account for a large share of intertemporal choice research (Cohen, Ericson, Laibson, and White 2016).

There are many anomalies that characterize measured RRR's. An influential review (Frederick, Loewenstein and O'Donoghue, 2002) documented numerous regularities uncovered in MEL experiments. For instance, as discussed above, the annualized RRR falls as the horizon, d , between t_1 and $t_2 = t_1 + d$ increases. There is also a robust magnitude effect (RRR's decrease as the magnitude of both x_1 and x_2 are proportionately scaled up). Moreover, as discussed in Section 2.5, there is *subadditivity* in MEL choices: individuals are more impatient over a delay if that delay is subdivided into smaller intervals (Read 2001; Read and Roelofsma 2003). Consider three delays: $t_1 < t_2 < t_3$, and two indifferences: $\{x_1, t_1\} \sim \{x_2, t_2\}$ and $\{x_2, t_2\} \sim \{x_3, t_3\}$. Transitivity implies that $\{x_1, t_1\} \sim \{x_3, t_3\}$, but subadditivity implies that the individual will strictly prefer $\{x_3, t_3\}$ to $\{x_1, t_1\}$. This complicates any easy aggregation of choices into a simple RRR.

Finally, the method used to elicit choices can affect measured preferences.²⁵ The literature has used a variety of methods, most commonly multiple price lists, matching/fill-in-the blank, or convex time budget tasks.

²⁴ We believe that the first implementation in humans was an unpublished working paper by Maital and Maital (1977). A related stream of work studied discounting in animal behavior (e.g. Chung and Herrnstein 1967, Ainslie 1975).

²⁵ Freeman, Manzini, Mariotti, and Mittone (2016) compare two "matching methods" (that of Becker DeGroot and Marshak 1964, henceforth BDM, and a sealed second price auction) with multiple price lists and find a significant

There is no single RRR that summarizes an individual's preference over money. However, we can summarize some of the quantitative stylized facts that have emerged from this literature. For payments now versus later, Thaler (1981) found median annualized RRRs of 20-30 percent for delays of 3-5 years, and annualized RRRs of 40-345 percent for delays of 1 month, depending on the magnitude of the rewards in his choice set. In the earliest paper to examine discounting in a representative sample (with incentivized choice), Harrison, Lau, and Williams (2002) conduct a field experiment in Denmark, and find an average annual RRR of 28% (with substantial individual heterogeneity), based on choices between an earlier payment in 1 month and later payments at delays of 6 to 36 months. Dohmen et al. (2010) study a representative sample of Germans who (repeatedly) chose between a payment now or a payment in 12 months, and report a median annualized RRR of about 30%.

An annual RRR of 30% is high (for instance, higher than almost all credit card APRs). However, note that these RRRs are estimated from choices with delays of months to years. Choices measured from delays of only days yield even higher annualized RRRs, as summarized by Frederick, Loewenstein and O'Donoghue (2002).

A second wave of experiments aims at measuring the discount function for a time-separable utility model. This typically requires accounting for potential curvature in the utility function.²⁶ The literature also typically assumes a "consume-on-receipt model" in which individuals consume the money on the date they receive it. The consume-on-receipt assumption is different from an optimization model in which individuals smooth consumption across periods. (See a detailed discussion below and in Cohen, Ericson, Laibson, and White 2016.)

Discount rates are typically lower once utility is adjusted for curvature. Andersen et al. (2008) measure utility curvature from choices over gambles (a "double multiple price list," once for time and one for risk). They estimate an average discount rate of about 10% per year for a representative sample of experimental participants in Denmark. Andreoni and Sprenger (2012) introduce a convex time budget method to account for utility curvature, and estimate an average annual discount rate of 25 to 35% for college students.²⁷

It is worth noting that utility curvature for small payments requires two assumptions: that payments are consumed on the day (or period, however defined) the payments are received (without crowd out), and that individuals have small scale risk aversion. The first assumption requires a strong degree of narrow bracketing, and is at odds with field evidence measuring marginal propensities to consume. The second assumption is empirically plausible, though it is at

difference between RRRs measured with the BDM mechanism and multiple price lists. Hardisty et al. (2013) compare matching with multiple price lists and find that matching has fewer demand characteristics. (Multiple price lists can anchor subjects or suggest a reasonable range of choices.)

²⁶ Abdellaoui, Attema, and Bleichrodt (2010), Olea and Strzalecki (2014), and Ericson and Noor (2015) provide alternative methods for measuring characteristics of discounting without measuring utility curvature. Noor (2009) highlights the importance of accounting for curvature and background consumption in reaching conclusions about the shape of the discount function.

²⁷ Andreoni, Kuhn, and Sprenger (2015) compare the performance of the convex time budget to the double multiple price list method (utility curvature measured from risky choices) and find that the convex time budget predicts better.

odds with the theoretical prediction of the expected utility model that utility should be approximately linear with respect to small gains or losses (Rabin 2000). Another line of work argues that attitudes over risk do not describe consumption utility curvature (see Andreoni and Sprenger 2012b, Abdellaoui et al. 2013, and Cheung 2016).

Another literature examines RRR for money outside the lab. This literature has some advantages over lab-based questions, but also raises methodological questions. One of the most influential papers was written by Hausman (1979), who examined the tradeoff between upfront prices when buying appliances v. the long-run operating costs (energy efficiency), and found implied annualized RRR's of about 20%. However, the tradeoffs may not have been salient to individuals (see a discussion in Allcott and Greenstone 2012).

There is also an ongoing debate about the RRR implied by individuals' willingness to pay for fuel efficiency: Dreyfus and Viscusi (1995) estimate an RRR of 11-17%; Allcott and Wozny (2014) estimate an RRR of about 15%. Busse, Knittel, and Zettelmeyer (2013) estimate RRR's that range from -6.2% to +20.9%, depending on their assumptions, and argue that the RRR is similar to the range of interest rates paid by borrowers.

In another classic paper, Warner and Pleeter (2001) examine the choice between an annuity and a lump sum payment. They show that most enrollees selected the lump-sum, implying a RRR that exceed 17%. This paper has been widely interpreted as evidence for under-annuitization. However, this estimate had important confounds: choosing the annuity came with additional requirements to be in the reserves; lump-sum recipients also received a variety of additional benefits not given to annuity recipients (see Simon, Warner, and Pleeter, 2014).

Similarly, Coile et al. (2002) find that individuals claim Social Security payments too early; under reasonable discount rates, most would be better off delaying claiming and effectively purchasing a larger annuity. This can be interpreted as evidence for a high RRR or may instead be linked to lack of knowledge about the rules and tradeoffs that exist in the Social Security system (Brown et al 2017). There is also a more general puzzle about low levels of annuitization (Benartzi, Previtro, and Thaler 2011; Beshears et al 2014).

Open Question #1: Why do individuals have such a high RRR for money if impatience is fundamentally about consumption flows rather than financial flows?

The money-earlier-or-later (MEL) literature raises fundamental methodological questions. A body of theoretical research argues against inferring discount rates from MEL studies because financial flows are fungible. Money flows are about financing, and distinct from the actual timing of consumption. Coller and Williams (1999) develop a model in which MEL choices reveal an individual's discount rate only if the revealed annualized RRR is strictly between her borrowing and lending rates.²⁸ In their framework, an individual's annualized RRR reflects either her financing costs/return or her discount rate. However, Cubitt and Read (2007) show that the censored data techniques of Coller and Williams are not applicable when agents have concave

²⁸ See also Chabris et al (2008) for another analysis of the confounds posed by intertemporal substitution, as well as other problems associated with the MEL paradigm.

utility functions and pick interior points of their choice sets: even when an observed RRR (imputed from binary choices) is between an individual's borrowing and lending rate, it will not reveal their discount rate but rather joint information about both the curvature of their utility function and their discount rate.²⁹

To utilize the MEL paradigm, the literature has often explicitly or implicitly assumed a consume-on-receipt model, in which individuals consume payments when they get them (e.g. Thaler 1981; Kirby and Herrnstein 1995; Loewenstein and Prelec 1992; Andersen et al 2008, Benhabib et al. 2010, Halevy 2014, 2015). Yet the estimated marginal propensity to consume out of financial receipts is much lower than one: people don't instantly consume everything they receive on the day of receipt. For instance, Johnson, Parker, Souleles (2006) examined spending responses to receipt of a tax rebate (\$300-\$600) that households knew would be coming and found that in the 3-month period after receiving the payment, households spend 20-40% of the payment on non-durables; Parker et al. (2013) found similar effects.

Evidence specific to MEL studies also provides evidence against the consume-on-receipt model. Reuben, Sapienza, and Zingales (2015) looked at the check cashing behavior for participants paid in MEL studies, and found that only about half of participants cashed their check within two weeks of receipt, even though they had chosen a smaller immediate payment than a 2-week delayed payment. However, note that consumption could occur without cashing the check.

One might rescue the consume-on-receipt assumption if experimental participants made choices *as if* they would consume on receipt, even if they do not. However, research we discuss in the next section suggests that individuals discount money differently than real consumption rewards (Augenblick, Niederle, and Sprenger 2015). Moreover, both Carvalho, Meier, and Wang (2016) and Dean and Sautmann (2018) show that MEL decisions were affected by income shocks, providing evidence against this narrow bracketing assumption (see also Cassidy 2017). Similarly, Krupka and Stephens (2013) show that changes in the inflation rate and in income are correlated with MEL decisions.

As a result, it remains an open question how to interpret behavior in MEL studies, and the extent to which it informs us about discount rates versus other factors such as heuristics, trust, risk, or probability weighting. (We discuss some of these alternative factors in Section 4.3.)

3.2. Preferences over Consumption Timing

Empirical Regularity #2: Individuals require a high rate of return (RRR) for real rewards. Choices over real rewards display robust present-focused preferences. When estimated with structural models, short-run discount rates are high.

²⁹ See Andreoni et al (2018) for evidence against arbitrage as an explanation for behavior on intertemporal choice tasks.

If discounting is over utility flows, a promising experimental design is to give individuals choices over time-yoked consumption, such as food or leisure—what we will term “real rewards.” The literature on real rewards and intertemporal choice was inaugurated by Mischel and Ebbesen (1970) and Mischel et al (1989) with the famous “marshmallow test” conducted with children. There were many different design variants, aimed at testing strategies that enabled self-control, but the different paradigms shared a common structure: participants were brought to a lab, and were told they could have one candy/cookie now or wait until the experimenter returned and have two. Marshmallows weren’t always used, but they have become the canonical example of how the experiment was run. Some subjects took the “marshmallow” immediately, others waited for some time before taking the sooner marshmallow, while still others waited for experimenter to return (typically 15 minutes) to receive both marshmallows. The marshmallow test is commonly assumed to measure individual differences in impatience, and some evidence shows that behavior in the test is predictive of later life outcomes.³⁰ However, Michel’s original intent was to investigate the effect of different strategies for implementing self-control (rather than persistent individual differences). Moreover, McGuire and Kable (2012, 2013) argue that participants who waited some time, then took a single marshmallow, might be better interpreted as learning about their environment (i.e. learning how long it would take for the experimenter to return) rather than failing to implement self-control.

Many studies show that RRR for real rewards are typically higher than for money: see e.g. Odum and Rainaud (2003), Odum, Baumann, and Rimington (2006), Estle et al. (2007), Lawyer et al. (2010), Tsukayama and Duckworth (2010), Reuben, Sapienza, and Zingales (2010), and Ubfal (2016). Some papers find extremely high RRR.³¹ McClure et al. (2007) employed time-dated juice rewards in a neuroimaging experiment in which participants were asked to choose one of two time-yoked water/juice squirts (subjects were thirsty, because they were denied fluid for three hours before the experiment and then fed salty snacks at the start of the experiment). They assume linear utility and estimate a quasi-hyperbolic discounting model that treats immediate juice as “now” but delays as small and 1 min in the future. They estimate $\beta = 0.52$ and δ close to 1.

Converting those RRR’s into discount rates requires the researcher to account for curvature in the utility function, where curvature is now not over money, but over the real reward itself. Augenblick, Niederle, and Sprenger (2015) show that participants in their study discount monetary rewards differently from real effort tasks. Estimating a quasi-hyperbolic discounting model separately for both money and effort, they find very little present-bias for money ($\beta = 0.97$) but more present-bias for effort: $\beta = 0.89$.

Testing for the existence of present-bias (or more generally present-focused preferences) does not require accounting for the curvature of the utility (however, quantifying present-bias does). Present-bias (or present-focused preferences) can be demonstrated when the *type* of real rewards

³⁰ A literature using small subgroup analyses found that performance in the marshmallow task predicted SAT scores years later (Shoda, Mischel, and Peake 1990). A recent conceptual replication (Watts, Duncan, and Quan 2018) found a correlation about half the size as the original correlation, and which was reduced by 2/3 when controls for childhood background were included. Benjamin et al (2018) report that the (childhood) marshmallow task does not predict mid-life measures of capital formation.

³¹ For more detail, see the discussion in Cohen, Ericson, Laibson, and White (2016).

chosen systematically differs when individuals chose for immediate consumption versus for future consumption. In this paradigm, participants make choices between “long-term gratification” and “short-term gratification” goods. For instance, Read and van Leeuwen (1998) examined the choice between healthy and unhealthy snacks. Across conditions, unhealthy snacks were chosen 51% of the time when offered for consumption one week in the future, while same unhealthy snacks were chosen 83% of the time when offered for immediate consumption. Similar results were shown by Read, Loewenstein, and Kalyanaraman (1999) in a paradigm in which participants choose between “high-brow” and “low-brow” movies. Studies like this make a conceptual distinction between goods that generate immediate utility costs and future utility benefits – what DellaVigna and Malmender (2004, 2006) call “investment goods” – and goods with the opposite time profile – or “leisure goods” (see the related discussion in Section 2, where we discuss the measurement of experienced utility, and Banerjee and Mullainathan 2010 on “temptation goods”).

Finally, in contrast to the experimental literature, studies using field choices have inferred discount rates explicitly for consumption. Laibson, Maxted, Repetto, and Tobacman (2017) use the method of simulated moments (MSM) to estimate a lifecycle consumption model with present bias, estimating $\beta = 0.5$ and $\delta = 0.987$ (see also Angeletos et al 2001). Other papers have used job search behavior to calibrate discounting models. DellaVigna and Paserman (2005) find β often near 0.9 (depending on assumptions). Paserman (2008) estimates β separately by income and finds that $\beta=0.4-0.5$ for low and moderate wage workers, but less present bias for high wage workers ($\beta=0.9$).

Food consumption provides another way to calibrate discounting models, and results are not consistent with exponential discounting. Using caloric intake data on food stamp recipients, Shapiro (2005) estimates high short-run discounting; imposing the exponential model, he estimates an unreasonable annual discount *factor* of 0.23—implying a discount rate of $-\ln(0.23)$ —but when estimating a quasi-hyperbolic model and assuming log-utility, he roughly calibrates $\beta = 0.96$. Mastrobuoni and Weinberg (2009) examine food consumption of liquidity constrained social security recipients and find $\beta = 0.91-0.94$.

Open Question #2: How substitutable is consumption (and effort) across time? How does the substitution of consumption affect measured discount rates?

Real rewards appear promising because they deliver incremental utility flows at a given time; moreover discounting of real rewards appears different than discounting of money receipt. Yet, in theory, consumption studies can introduce similar confounds: individuals can change consumption in other areas of life to offset experimentally induced consumption. For instance, consider an individual choosing how much to work (as opposed to relax) in a lab session today v. next week. If the individual works more in the lab right now, they could offset that by relaxing more after they leave the lab. Moreover, present choices may appear different from future choices: if an individual’s plans for today are largely fixed but future plans are flexible, effort in the lab today will be incremental effort for the day, while planned effort for next week can be offset by reducing effort outside the lab. If, for instance, utility is aggregated at the day level and can easily be adjusted outside the lab, the utility experience in a given hour in the lab need not impact utility that day.

As intertemporal choice research increasingly moves toward using real consumption rewards, estimates of how substitutable utility is across time and structural models that account for offsetting behavior are needed.

3.3. Preference Reversals

Empirical Regularity #3: Individuals exhibit (present-focused) preference reversals.

One of the key anomalies of the intertemporal choice literature is the concept of preference reversals. Indeed, in our view, the unifying theme of the (behavioral) intertemporal choice literature is the concept of preference-focused preferences (introduced in Section 2), which boils down to a preference reversal. Recall our earlier definition: *Present-focused preferences exist if agents are relatively more likely to currently choose actions that generate instant gratification, then they are to currently choose the same (binding) actions for future periods.* We hasten to remind the reader that preference reversals are not the same as dynamic inconsistency in preferences, a point we will return to below.

Read and Van Leeuwen (1998) present a canonical example of a preference reversal. In their study experimental participants are asked to choose among six snacks (two healthy and four unhealthy)³² to be delivered after seven days. We'll refer to the asking date as day 0 and the delivery date as day 7. On day 0, the participants are not told that they will have an opportunity to revisit this choice on day 7, an important design decision, that we will discuss momentarily. Read and Van Leeuwen create a 2×2 design, where they vary the degree of satiation at day 0 ("after lunch time," when they are presumably satiated, or "in the late afternoon, around 4:30 or 5:00," when they are presumably not satiated) and the degree of satiation at day 7. Experimental participants are told on day 0 when the snack will be delivered on day 7. In other words, they know at day 0 whether they will receive the day-7 snack after lunch time or in the late afternoon.

In the always satiated condition (asked when satiated and delivered when satiated), 26% of participants choose an unhealthy snack at date 0 (to be delivered on day 7), and 70% of participants choose an unhealthy snack on day 7. In the always non-satiated condition (asked when non-satiated and delivered when non-satiated), 78% of participants choose an unhealthy snack at date 0 (to be delivered on day 7), and 92% of participants choose an unhealthy snack on day 7. (There is a ceiling effect in the always non-satiated condition.)

In both the always satiated condition and the always non-satiated condition there is a switch on day 7 towards snacks that offer greater instant gratification (e.g., a Snickers bar vs. an apple). This is the essence of present-focused preferences. When people choose the snack to be eaten "now" they are relatively more likely to choose snacks that offer instant gratification than when they choose (what they think will be a binding) choice 7 days in advance. Many of the models

³² The six snacks, rank-ordered from most to least healthy are apples (1.33), bananas (1.67), crisps (3.91), borrelnoten (4.08), Mars bars (4.67), and Snickers bars (5.33). The numbers in parentheses are the average rank-ordering.

discussed in Section 2 would predict this type of “preference reversal,” although they offer different explanations for the reversal.

Many related studies have documented preference reversals, including Read, Loewenstein and Kalyanaraman (1999), DellaVigna and Malmendier (2004, 2006), Badger et al (2007), Milkman, Rogers, and Bazerman (2009), Augenblick, Niederle, and Sprenger (2015), Sadoff, Samek, and Sprenger (2015), Kuchler and Pagel (2018), and Fedyk (2018).³³ It is also important to note that preference reversals come in three broad methodological categories. First, strong form preference reversals occur when the same decision is made by the same person at an early date and then again at a later date (as in Read and Van Leeuwen 1998). Semi-strong form preference reversals occur when the same person makes a decision now and also makes a decision for the future. Weak form preference reversals occur when different people make the now decision and the later decision.

Open Question #3: What are the key mechanisms that cause preference reversals?

Preference reversals are consistent with *both* dynamically *inconsistent* preferences *and* dynamically consistent preferences. For example, present-bias generates preference reversals of the type that we have just discussed due to dynamically *inconsistent* preferences (see Section 2.1). Present-biased agents heavily weight immediate rewards (which are not β -discounted) relative to delayed rewards (which are β -discounted), leading present-biased agents to more frequently choose candy (and other unhealthy snacks) when they are available immediately, compared to when they are available at a delay. Read and Van Leeuwen (1998) interpret their results as evidence for present-biased preferences, but this is not the only way to understand their findings.

Unitary-self models (which feature dynamically consistent preferences), also predict preference reversals. Here the mechanism is the temptation effects that are assumed to arise when decisions are being made in the moment (i.e., for eating candy now) but do not arise when decisions are being made for the future. In other words, candy now is tempting, but candy next week is not tempting. Accordingly, an agent will more readily choose candy now than candy later. All of this can be explained with a model in which agents are dynamically consistent. Specifically, the agent says that she would like to choose fruit later *if* she can commit herself. However, she does not prefer to choose fruit later if she can't commit herself, because she anticipates that the temptation effects (that will arise later if she hasn't pre-committed) will make choosing fruit later suboptimal. In this sense, she is dynamically consistent. On day 0, she only prefers fruit later if she can pre-commit herself to choose fruit later (i.e., if she can restrict her choice set). On day 0, she prefers candy later if she can't restrict her day-7 choice set.³⁴

Of course, (dynamically inconsistent) present-bias and (dynamically consistent) temptation effects are mutually compatible. In other words, they are likely both present in human motivation. One open challenge is to identify and measure their quantitative contributions. A related open challenge is to understand the role that naiveté plays in all of these intertemporal

³³ Preference reversals also emerge in other domains. For example, Andreoni, Aydin, Barton, Bernheim, and Naecker (2016) document preference reversals in social preferences.

³⁴ See Sadoff, Samek, and Sprenger (2015) for related evidence from a field experiment.

choice models. For example, a naïve agent with temptation preferences may not appreciate how temptation will affect her later choices, leading the agent to look nearly indistinguishable from a naïve present-biased agent. Accordingly, it is important to understand not only the underlying mix of preferences that agents have, but also to understand their meta-awareness of these preferences.

3.4. Procrastination

Empirical Regularity #4: We procrastinate and deadlines can help.

While many people agree that procrastination is widespread, we do not have precise, commonly shared definition of procrastination. In our discussion here, we will define someone as procrastinating if, when completing a costly task, there is delay that appears suboptimal from their *own* perspective. Procrastination, is by our definition, a problem, and can result from different models.³⁵

Anecdotal examples of procrastination abound, perhaps most famously in Akerlof's (1991) discussion of how he continually put off returning a box to a friend—and showed how small mistakes could lead to large welfare losses: each day, the loss from delaying one more day was small, but the overall cost of delaying versus sending at the start was quite large.

The theoretical analysis of procrastination through the lens of present-bias was developed in a series of papers by O'Donoghue and Rabin (1999a, 1999b, 2001). For naïve individuals completing a single task, present-bias quite intuitively leads to excess (and unpredicted) delay; but behavior of present-biased individuals can become quite complex for sophisticates, or when there are multiple tasks from which to choose. For instance, there are special circumstances in which sophisticates may sometimes complete a task *earlier* than a time-consistent individual because they recognize if they delay today they will delay even more in the future; as a result, present-bias needn't always lead to procrastination. However, in most settings, agents with present bias (i.e., $\beta < 1$, holding all else equal), will procrastinate more than agents without present bias, whether the present-biased agents are naïve or sophisticated.

Nor does every behavior that looks like problematic delay qualify as procrastination. For instance, it is widely observed that people—especially students—wait until right before a deadline to act, and this is often interpreted as procrastination. Yet this behavior can emerge quite naturally if there is little value to completing the task in advance of the deadline (see Fisher 2001's theoretical analysis). If individuals did not describe waiting until the deadline to be inefficient, this wouldn't qualify as procrastination in our definition.

Much of the systematic evidence on the extent of procrastination comes from the laboratory (Augenblick, Niederle, and Sprenger 2015) or from field experiments with students (e.g., Fedyk 2018). There is also an older literature in education, not well-known by economists, examining

³⁵ Alternatively, one can view procrastination through the lens of a model and define it more precisely. For instance, through the lens of present-bias, procrastination can be defined as additional delay in completing a costly task relative to the benchmark of a time-consistent model with correct beliefs (e.g. as in Ericson 2017).

procrastination in students. For instance, Solomon and Rothblum (1984) found many students self-report “procrastinating” in preparing for exams and that they also report wanting to reduce procrastination, though it is unclear how subjects defined procrastination.

Procrastination is often closely connected to naiveté (but naiveté is not necessary for procrastination – see Carroll et al 2009). DellaVigna and Malmendier (2004, 2006) have a pair of papers that examine procrastination behavior in the context of gym attendance; we discuss the role of naiveté in these papers in the next section. DellaVigna and Malmendier (2006) show that individuals take “too long” to cancel their gym contract after their last attendance (an average of about 2 months). Calibrations suggest that a sophisticated person would cancel in a few days, while naiveté can explain the much longer observed delay. DellaVigna and Malmendier (2004) show that firms will design contracts to take advantage of this procrastination on cancellation behavior, creating contracts with automatic renewal and switching and cancellation costs.

Other research identifies behavior that is likely procrastination and shows it is connected to other behaviors that result from present-focus. For instance, Brown and Previere (2016) show the people who wait until the last day of health plan open enrollment to enroll (“procrastinators”) take longer to sign up for 401(k) plans, contribute less to those plans, and are more likely to take the default fund allocation in the 401(k). Similarly, Reuben, Sapienza, and Zingales (2015) measure how quickly individuals complete various tasks (e.g. application to university and mandatory surveys); procrastinators on those measures take longer to cash their payment check.

Deadlines are a type of commitment device that may reduce procrastination. Ariely and Wertenbroch (2002) provide influential evidence on deadlines that is often interpreted as evidence about procrastination. They assigned one section of students to fixed, evenly spaced deadlines for class assignments, while students in the other section got to choose their own deadlines at the beginning of the semester. One key fact is that when given the choice, about two-thirds of their study participants do self-impose binding deadlines (with financial penalties), suggesting some sophistication about a self-control problem. However, these self-imposed deadlines were less binding, on average than the evenly spaced deadlines. Second, performance on assignments was lower in the self-chosen deadline condition than in the evenly spaced deadlines. They show similar results in a lab-based proofreading task. The paper suggests the interpretation that students do not set self-imposed deadlines “optimally,” from perspective of maximizing class performance. However, it is possible that students in the self-chosen deadline condition could be trading off preferences for flexibility in other areas of life against performance in this class; in this view, self-chosen deadlines could lead to lower class performance but higher individual welfare.

However, it is not always the case that interim deadlines increase performance; there is a tradeoff between overcoming procrastination and maintaining flexibility. Burger, Charness, and Lynham (2011) conduct deadline experiments for subjects completing online tasks. They found that interim deadlines lowered performance (as measured by completion rates), in contrast to Ariely and Wertenbroch’s (2002) results. Similarly, Bisin and Hyndman (2018) examine an experiment in which participants have to complete tasks online. While they find robust demand for self-imposed deadlines, they also find that deadlines did not increase completion rates. Bisin and

Hyndman estimate a structural model of present-bias and conclude that present-bias is widespread.

Evidence on procrastination and deadlines outside of students and lab experiment is more limited (however, see, Kaur, Kremer, and Mullainathan 2010, 2015; Duflo, Kremer and Robinson, 2011). Frakes and Wasserman (2017) provide intriguing evidence on procrastination in patent examiners. They find that work is “end loaded” right before deadlines and argue that non-procrastinating examiners would spread work evenly over time, while a procrastinator who wants to make their target would cluster work at the end of each quota period. They provide evidence that end loaded work is of lower quality. Moreover, it is likely that telecommuting (as opposed to working in the office) could exacerbate procrastination; indeed, they find that end-loading increases when a telecommuting program is introduced.

Explanations for Procrastination

Procrastination is often explained by time inconsistent preferences. However, limited attention or memory can also explain excess delay, and can interact with present-bias in counter-intuitive ways (see Ericson 2017). For example, Choi, Laibson, Madrian, and Metrick (2002) examine procrastination behavior in the context of retirement savings for white-collar workers. Out of every 100 workers, 68 say that their savings rate is too low. When asked about future intentions, 24 of those 68 say they plan to increase their savings rate in the next two months. Yet data on contribution rates reveal that only 3 of those 24 actually do so (over the next four months). Similarly, at a financial education seminar, virtually all non-participating employees said they planned to join the retirement savings plan, but only 14% followed through. A combination of delay due to present-bias and then forgetting may explain this, as employees typically either make changes almost immediately after their financial education seminar or not at all.

Another explanation for procrastination is the planning fallacy. Kahneman and Tversky (1979) identified the concept of the *planning fallacy*: people typically underestimate how long it will take to complete a task. The planning fallacy could be in part due to mispredicting future present focus; the phenomenon is typically closely linked to incorrect beliefs about how much work will be required. However, evidence on planning fallacy suggests that misprediction has an important component apart from any issues of present-focus (e.g., a military contractor mispredicting the time to delivery on a major weapons project).

Open Question #4: Why do people underestimate their own procrastination? What are the relative roles of naiveté about present-focus versus other explanations for underestimating procrastination, such as overconfidence about the effort required and limited memory?

Evidence on forecasts of task completion (e.g. Augenblick and Rabin 2018, Fedyk 2018) show the people underestimate how long it will take to complete a task; estimated models of present-bias also often show substantial naiveté (see following section).

One open question is why people are naïve about their present-focus, which we discuss more extensively in the next section. However, there are at least two other explanations for

underestimating procrastination that do not rely on present-focus: overestimating memory (attention) and underestimating how much work a task will be.

Individuals might not complete a task because they forget to do it for some period of time. The issue of limited memory is taken up in a pair of papers by Ericson (2011, 2017). Ericson (2011) finds experimental evidence that individuals are overconfident about the probability they will remember to complete a task in the future. Ericson (2017) models the interaction between limited memory and present-bias and shows that structural models can produce biased estimates of present-bias if they incorrectly assume complete memory.

Empirically, lengthening deadlines can lower eventual task completion rates (Shu and Gneezy 2010; Taubinsky 2014). Having more time to complete a task should not lower completion rates in a model of present-bias but can lower completion rates with limited memory (see Taubinsky 2014, Ericson 2017). This suggests that limited memory or inattention is important for explaining procrastination, yet we do not have calibrated estimates of this mechanism.

3.5. Naiveté

Empirical Regularity #5: People partially but do not fully anticipate their extent of future present-focus.

Do individuals correctly anticipate the extent of their future present-focus? That is are they sophisticated, or naïve? A body of evidence indicates that people are at least *partially* naïve about their present-focus: they may recognize they will in the future act in a present-focused manner, but do not recognize the extent of this present-focus. Some evidence suggests that individuals are fully naïve (they think they will act like a time-consistent exponential discounter in the future).

We focus on naiveté about present-focus, which is a subset of incorrect beliefs. As a concept it is distinct from overconfidence about ability, over-optimism about future challenges, or projection bias. Nonetheless, it can be difficult in practice to distinguish among competing types of incorrect beliefs. For instance, not saving much for the future today could result from expectations that you will save more tomorrow, or that you will have a much higher income tomorrow. Similarly, persistent failure to complete a task could arise from present bias and naiveté (O'Donoghue and Rabin 1999a, 1999b) or from continually and incorrectly believing that the cost of completing the task will be lower tomorrow than it is today. Moreover, Ericson (2017) shows that overconfidence about memory/inattention can lead to a failure to act and may be mistaken for present-bias.

Early evidence on naiveté comes from DellaVigna and Malmendier (2006), who show a pattern of naiveté in the choice of gym contracts. First, individuals too often choose a contract with a flat fee (at an effective rate of \$17 per expected visit), even though the pay per visit is cheaper (\$10 per visit fee). Moreover, users in the monthly contract are likely to stay enrolled beyond one year than users committing for a year. This is surprising because monthly members pay higher fees for the option to cancel each month, and suggests they over-estimated the probability they would actually follow through on canceling.

While DellaVigna and Malmendier (2006) rely on contract choice to identify misprediction of future gym attendance, Acland and Levy (2015) directly elicit beliefs about future attendance and find that individuals over-predict their future gym attendance by 1-2 visits per week. (Moreover, participants in their study also failed to anticipate the extent to which incentives would increase their future behavior, which Acland and Levy attribute to projection bias.)

Ariely and Wertenbroch (2002)'s finding that equally spaced deadlines lead to better performance than self-chosen deadlines may be evidence of partial naiveté. However, as we suggested in Section 3.4, participants in their study may have valued the flexibility of the later, self-chosen deadlines.

Stronger evidence on incorrect beliefs in the classroom setting comes from Wong (2008), who asks students when it would be ideal to start preparing for the final and when they predict they will start preparing for the final. On average, students predicted they would start about 2.5 days later than ideal, which Wong interprets as awareness of a self-control problem. However, subsequent data on when students actually start preparing for the exam indicate that they started preparing an average of about 2 days later than predicted; Wong classifies more than half the sample as either fully or partially naïve and only a quarter as sophisticated (predictions within +/- 1 day of actual). Wong's results could indicate naiveté about present-bias but could also be consistent with overconfidence about future free time or about remembering to prepare.

Augenblick and Rabin (2018) examine present-bias and sophistication/naiveté about present-bias in a real effort task. They estimate a β of about 0.8 but estimate a $\hat{\beta}$ near 1 (individual's implied perception of the β that will determine their choices in the future, see Section 2.1). They cannot reject the null hypothesis of no perceived present bias and can reject hypothesis that individuals have an accurate perception of present bias.

Whether individuals are sophisticated or naïve about their present-focus matters for the contracts that individuals will choose in equilibrium as well as for welfare. For instance, firms will design contracts with backloaded fees if individuals are naïve about the probability they will procrastinate on taking costly action to cancel a contract (e.g. gym membership) in the future (DellaVigna and Malmendier 2004). Heidhues and Koszegi (2010) show that even a small amount of naiveté in credit contract choices will lead individuals to have discontinuously lower welfare than a sophisticated individual: firms design credit contracts with large penalties for deferring repayment and naïve individuals choose these contracts because they underestimate the probability they will pay these penalties.

Evidence suggests that people hold more accurate beliefs others' procrastination than their own. For instance, Fedyk (2018) estimates a structural model of present-bias on participants completing a real-effort task. While individuals' forecast of their own behavior implies a substantial degree of naiveté, their forecasts for others were much more accurate.

Using field data on consumption, Kuchler and Pagel (2018) provide a method to classify people as sophisticated or naïve based on how an individual's level of financial resources affects the sensitivity of their consumption to paycheck receipt. Their model shows that sophisticates, but

not naïfs, will act more patiently when there are more resources available. They examine credit card payments and compare self-set debt paydown plans to actual behavior. While many users fail to follow through, individuals their method identifies as sophisticates are in fact more likely to stick to their plan.

Open Questions #5: Why don't people learn and anticipate their present-focused behavior? What can help people correctly anticipate their future present-focused behavior?

There are a number of open questions about naïveté. Given that individuals repeatedly make consumption decisions, it is perhaps surprising that they do not learn (more quickly) about their miscalibrated beliefs. One hypothesis is that individuals do learn, but learn in narrow, domain-specific contexts. Therefore, they do not learn generally about the fact they make present-focused decisions. Under this hypothesis, individuals might learn that they will procrastinate on preparing for exams, but when they encounter a new decision (e.g. job search), they don't realize they will procrastinate. Yet we see naïveté even in contexts where individuals have experience (e.g. students in the study of Wong 2008), so they may not even learn in specific domains. It is likely that wishful thinking plays some role here. People would like to believe in a rosy future, and naïve beliefs help to make that possible. This would explain why people are asymmetrically naïve about their own future behavior, and not about the behavior of others (see Fedyk 2018).

3.6. The Effect of Transactions Costs

Empirical Regularity #6: Small changes in transactions costs often have large effects on behavior.

The behavioral economics literature has emphasized that small changes in transactions costs or choice architecture sometimes have anomalously large changes on behavior (for reviews see Thaler and Sunstein 2008 or Beshears et al 2018, which is Chapter 3 of Volume 1 of this Handbook). Default settings (Madrian and Shea, 2001; Johnson and Goldstein, 2003; Beshears et al 2009; Choi, Laibson, Madrian, 2009) are posed as the canonical example of such effects. For example, setting a default of non-enrollment in a 401(k) plan leads to a participation rate of approximately 40% after one year of tenure as an employee. By contrast, setting a default of enrollment (and letting employees opt out if they prefer), leads to a participation rate of approximately 90% after one year of tenure. Results like this have been reproduced in dozens of studies (for a meta-analysis see Jachimowicz et al 2018).³⁶

Other examples of highly impactful transaction cost manipulations abound. For example, offering parents of a high school senior immediate assistance and a streamlined process to complete the Free Application for Federal Student Aid, raised the likelihood that their child completed (at least) two years of college from 28% to 36% (Bettinger, Long, Oreopoulos, Sanbonmatsu, 2012). Enabling individuals to almost effortlessly opt into a retirement savings

³⁶ Note however that defaults have complex distal effects that often dilute their impact (Abadie and Gay 2006; Beshears et al 2018; Choukhmane 2018).

plan at a pre-selected contribution rate and asset allocation, increased enrollment rates by 10 to 20 percentage points (Beshears et al 2013). Enabling students to send an extra free ACT report to a college (rather than charging \$6 for that marginal report), led ACT-takers to send more reports and applications relative to SAT-takers, widening the range of colleges they sent scores to and improving admissions to selective colleges (Pallais 2015). Related effects have been observed in take-up of social benefits (Currie 2004; Aizer 2007; Bhargava and Manoli 2015) and medical adherence (Thornton 2008; Banerjee et al. 2010; Milkman et al. 2011, 2013).

Open Question #6: Is the large effect of small transactions costs on behavior primarily related to present focus, or is it some other channel?

The outsized impact of small incentives, including disincentives like transactions costs, is likely due to many different mechanisms. Present bias (and other closely related present-focused mechanisms) have been used to explain why agents are particularly responsive to immediate transactions costs and immediate rewards for task completion (e.g., Carroll et al 2009). These immediate transaction costs can be particularly impactful because they lead agents who are partially or fully naïve to postpone doing the task because they falsely believe that they are going to complete the task in the near future (e.g., Akerlof 1991; O'Donoghue and Rabin 1999). This catalyzes a long (or even infinite) cycle of self-defeating procrastination. Related empirical evidence points to present bias as a key mechanism that explains why transactions costs (and defaults specifically) are so effective in influencing behavior (e.g., Brown and Previtro 2016; Goda et al 2018; Blumenstock, Callen, and Ghani, 2018).

However, there are other mechanisms that are also probably important. For example, status quo bias (Samuelson and Zeckhauser, 1988) is likely driven by many complementary channels (not just present-focused preferences), including confusion, anchoring effects (e.g., Bernheim, Fradkin, and Popov, 2015), endorsement effects (e.g., Benartzi 2001; Brown, Liang, and Weisbenner 2007; Beshears et al. 2008), and inattention/memory (e.g., Gabaix 2014, Ericson 2017).

As nudges become increasingly popular in the policy context, it will be useful to identify the rich set of mechanisms that cause these nudges to be effective. Unpacking the many complementary mechanisms and identifying their quantitative contribution will be important as we push this research program forward.

3.7. Lack of liquidity on household balance sheets

Empirical Regularity #7: Households tend to hold very low levels of liquid assets.

Households around the world tend to hold very low levels of *liquid* assets, despite the fact that they tend to hold relatively large stocks of *illiquid* assets. These facts are compiled in Chapter 3 of Volume 1 of this Handbook (Beshears et al 2018, Table 1 and Appendix A), and we quickly touch on them here as well. These results are derived from the 2016 (U.S.) Survey of Consumer Finances.

The median value of net liquid assets starts at \$1,000 for households in the 21-30 age category and rises to \$6,719 for households in the 61-70 age category. (Other percentiles are presented in Chapter 3 of Volume 1.) Accordingly, even households on the brink of retirement have little liquid net worth.

On the other hand, total net worth (which includes illiquid assets nets out corresponding collateralized liabilities, and *excludes* Social Security claims, which are involuntarily accumulated) shows robust growth over the life course. The median value of total net worth starts at \$7,611 for households in the 21-30 age category and rises to \$209,227 in the 61-70 age category.

These results jointly imply that the typical U.S. household is living without a large buffer of liquid assets and doing almost all of its voluntary wealth accumulation in illiquid assets.

The financial health of U.S. households appears more strained when debt is broken out of the balance sheet. In any given month, half of households are borrowing on at least one of their credit cards (i.e., not making their full payment when they pay their card – see Zinman 2015 and Laibson, Maxted, Repetto, and Tobacman 2017). Finally, in any given year 12 million U.S. households (5% of the adult population) take out a payday loan (Pew Charitable Trusts, 2012). About one million households file for personal bankruptcy each year and one in ten households has filed for bankruptcy at some point in their lives (Stavins, 2000).

As a corollary to these low levels of liquid wealth (and high levels of borrowing), households also tend to have anomalously high propensities to consume out of changes in liquid wealth, whether that change in liquidity is anticipated or unanticipated (e.g., Shea 1995; Gross and Souleles 2002; Stephens 2003; Shapiro and Slemrod 2003; Shapiro 2005; Mastrobuoni and Weinberg 2009; Parker et al. 2013; Broda and Parker 2014; Ganong and Noel 2017; Gross, Notowidigdo, Wang 2016).

Aguiar and Hurst (2005) offer a countervailing view, emphasizing that the simultaneous fall in income and consumption expenditure at retirement does not coincide with a fall in caloric consumption. They argue that the systemic drop in consumption *expenditure* is *not* anomalous because households maintain a smooth consumption of calories between the pre- and post-retirement periods (by shopping more frugally and doing more home production of meals). However, Stephens and Toohey (2018) are unable to replicate this caloric-smoothing result using new data and instead find that caloric intake falls as households enter retirement (mirroring the fall in consumption expenditure).

Open Question #7: Why do households have such low levels of liquid net wealth, relatively high levels of illiquid net wealth, and high marginal propensities to consume out of liquid wealth changes?

These stylized facts are particularly hard to explain because these empirical regularities tend to occur in the same households. For example, households that borrow on their credit cards tend to also accumulate large stocks of illiquid wealth (see Laibson, Maxted, Repetto and Tobacman

2017). Accordingly, household heterogeneity can offer only a partial explanation of these regularities.

In the behavioral economics literature, the principle explanation for these balance sheet and consumption facts has been present bias (Laibson, 1997, Angeletos et al 2001; Skiba and Tobacman, 2008; Laibson, Maxted, Repetto, and Tobacman, 2017). Households with present bias have a hard time holding on to liquid assets, which tend to get quickly spent on consumption expenditures. Indeed, households with present bias (e.g., $\beta < 0.8$) will tend to borrow on credit cards. However, present biased households will be willing to invest in illiquid assets (including durables) with modest rates of return, because the household's *long-run* discount rate is only $\ln \delta \cong 1 - \delta$, which is typically estimated to be below the long-run real (after-tax, risk adjusted) return that households can obtain on assets like housing or (matched) 401(k) contributions. Laibson, Maxted, Repetto and Tobacman (2017) use the method of simulated moments to estimate preference parameters in a lifecycle consumption model. Using data on voluntary wealth formation and credit card borrowing to identify preference parameters, they estimate $\beta=0.51$, $\delta=0.99$, and a coefficient of relative risk aversion of 1.3.

Models outside the behavioral literature have taken a different approach to explaining low levels of liquidity, high levels of borrowing, high levels of illiquid wealth formation and the high marginal propensity to consume out of changes in liquidity. For example, Kaplan and Violante (2014) assume that the real, after-tax return on liquid wealth is -1.48%, the real, after-tax interest rate on credit card borrowing is 6%, and the real, after-tax and risk-adjusted return on illiquid assets is 6.29%, generating a high level of credit card borrowing, little liquid wealth buffering, and aggressive investment in illiquid assets (see also and Kaplan, Violante and Weidner, 2014).

Future work will continue to explore the mechanisms that explain the balance sheet behavior and consumption dynamics of households.

3.8. Commitment

Empirical Regularity #8: People sometimes demand commitment devices.

Many of the theories that generate present-focused preferences predict that, under the right circumstances, agents will choose commitment 'for its own sake' (see Table 1 in Section 2.7). By this we mean that agents will prefer to restrict their choice set holding all else equal, which we refer to as a pure commitment. For present-biased agents, such a restriction prevents them from making a suboptimal choice in the future (e.g., a pre-paid or shame-inducing gym trainer forces one to exercise; likewise an irreversible temporary web-or phone-blocker like [Freedom](#) keeps one on task). For agents with temptation (see Section 2.2), commitment removes temptation costs, which can be beneficial *even* if the tempting good is not chosen when it is available.

Agents sometimes do choose *pure* commitments that are not confounded by tied inducements, like financial incentives,³⁷ but such pure commitments are rarely observed in markets (Laibson, 2015). When trying to identify pure commitments in the field, only a handful come to mind, such as web blockers, which do not seem to be used widely, even in developed economies.³⁸ We suspect that personal trainers are also hired mostly for their commitment effects, though this is not an example of a pure commitment.

Most practical “commitments” seem to be corrupted by impurities. For example, standard amortization mortgages (with fixed payments that end when the principle is repaid) are sometimes advocated in the financial advice literature as a forcing mechanism that helps people with self-control problems systematically shed their housing debt. But mortgages are also a tax-advantaged financing vehicle for home buying. So it’s not clear why people take them up: to solve a self-control problem or, simply, to get tax advantaged financing. Perhaps the lack of popularity of interest-only mortgages provides a hint that amortization mortgages deliver some special benefits, but such an inference is somewhat speculative.

Some other *apparent* commitments may simply arise from naiveté. For example, Wertenbroch (1998) documents that consumers prefer to buy vice goods in small packages and virtue goods in large (volume discount) packages, which is reflected in firms’ pricing decisions. These patterns could arise because of commitment *or* because consumers *naively* anticipate that they won’t want to eat more than a small serving of ice cream and accordingly buy ice cream in packages that forego volume discounts. If you think you are only going to eat ice cream today (because you believe you will prefer to eat healthfully tomorrow), then you should buy only a small container of ice cream. Hence, small packaging could be a sign of commitment or it could be a sign of naiveté.

The commitment evidence from experiments is less muddy. In lab studies and field experiments researchers have been able to systematically observe pure commitment behavior (e.g., see Ariely and Wertenbroch 2002; Ashraf, Karlan, and Yang 2006; Bryan, Karlan and Nelson, 2010; Houser, Schunk, Winter, and Xiao, 2018; Gine, Karlan, and Zinman 2010; Kaur, Kremer, and Mullainathan 2010, 2015; Royer, Stehr, and Sydnor, 2015; Bisin and Hyndman 2018; Augenblick, Niederle, and Sprenger 2015; Alsan et al, 2017; Casaburi and Macchiavello forthcoming; Schilbach forthcoming; Beshears et al, 2018). For example, Schilbach (forthcoming), implements a field study in which the experimental participants are rickshaw drivers in India, a population with a high level of alcohol consumption. Schilbach finds economically (and statistically) significant demand for a pure commitment device. About a third of study subjects choose to forgo about 10% of their daily incomes for a week to receive monetary payments conditional on zero blood alcohol test that would otherwise be received unconditionally.

³⁷ For example, a 401(k) savings plan is an *impure* commitment, because it is tied to financial incentives like matching dollars and favorable tax treatment, so it is not clear whether people invest money in 401(k)’s for their commitment properties or despite them.

³⁸ Other exotic examples include a series of [alarm clocks](#) that jump off the night table and run around the bedroom until the groggy sleeper chases them down. Such alarm clocks do exist, but they appear to have a limited market. Likewise, the commitment website [stickK](#) has not gone viral.

Open Question #8: Why is pure commitment so rare in markets and why is willingness to pay for commitment usually so low (even in lab experiments designed to measure the taste for commitment)?

As noted above, pure commitment is rarely observed in organic ('real-world') markets that have not been designed by behavioral economists (e.g., Bernheim, Meer, and Novarro, 2016). While pure commitment is routinely observed in lab and field experiments (see the list above), willingness to pay for such commitment is typically zero (e.g., Augenblick, Niederle and Sprenger 2015). In other words, the typical experimental participant is not willing to *pay* for the privilege of tying their hands (by committing to a particular outcome that is strictly dominated by what they could have achieved without commitment). Only a handful of studies reveal a substantial willingness to pay for commitment: e.g., Casaburi and Macchiavello (forthcoming), and Beshears et al (2018), report that an overwhelming majority of participants are willing to pay for commitment, and Schilbach (forthcoming) finds that a third of subjects are willing to pay 10% of their daily income for commitment.

There are several hypotheses for the overall pattern of weak demand for commitment (see Laibson 2015, 2018). First, agents may be fully or partially naïve. Second, commitment may be costly to generate (e.g., personal trainers). Third, extrinsic uncertainty may make commitment suboptimal (if the commitment technology can't condition on all states of nature). Fourth, many types of commitment may already be embedded in other institutions that serve multiple purposes. For example, the 401(k) system provides both a commitment technology – penalties for early withdrawal – and a saving subsidy, crowding out incipient demand for a pure commitment technology (e.g., a savings account that is illiquid and offers no saving subsidy). Many social institutions may have this hybrid property of providing commitment and other services. For example, many large employers provide multiple forcing mechanisms that counteract workers' self-control problems (e.g., deadlines, progress reports, retirement savings programs, etc. – see Laibson 2018). Finally, self control may be achieved through internal punishments instead of external commitment technologies (see Bernheim, Ray, and Yeltekin 2015).

3.9. Paternalistic Policy and Welfare

Empirical Regularity #9: Many public policies seem paternalistic and aimed at addressing present-focused behavior.

Many government policies seem rooted in a paternalistic desire to at least partially offset present-focused behavior. The most explicit is provision for retirement: in the U.S., social security benefits offer protection to myopic or present-focused individuals who do not save for retirement. Indeed, as far back as in Feldstein (1985), setting the optimal level of social security benefits was discussed as a tradeoff of benefits for myopic individuals against costs from distorting private saving and other economic behavior. Illiquid defined contribution retirement savings accounts also help provide a commitment device for savers. Laibson, Repetto, and Tobacman et al. (1998) show that this commitment can be appealing for present-biased individuals. (Note that U.S. accounts are more liquid than in some other developed countries; see

Beshears et al. (2015) for a discussion.) Several recent papers have acquired that illiquid accounts and compulsory savings are an optimal response if agents have present-biased preferences (e.g., Amador, Werning, and Angeletos 2006; Moser and Olea de Souza e Silva 2017; Beshears et al 2018).

Compulsory education laws are also naturally—and often explicitly—explained as a counter to present-focused decisions of children (see Lavecchia, Liu, and Oreopoulos 2014). Restrictions on payday borrowing as well as caps on interest rates (usury laws) can be motivated in part by an intent to counter present-focused behavior (see e.g. Skiba 2012 on what restrictions on payday lending can be rationalized by present-focused behavior and what cannot).

Paternalism can interact with present-focused behavior to lead to somewhat surprising predictions: for instance, Lockwood (2016) shows that with present-bias can lead to optimal tax policy to have negative marginal tax rates at low incomes, similar to the design of the Earned Income Tax Credit in the U.S.

The existence of paternalistic policy plus naiveté about present-focus may at first seem a puzzle, but recall that individuals are better at predicting the present-focus of others than themselves, and these laws can be interpreted as paternalistic control of others, rather than a preference for commitment (Fedyk 2018; also see Albrecht et al 2011). Moreover, Laibson (2018) shows that naïve individuals can use model-free forecasts to choose contexts (e.g. firms/employers) that have a reputation for generating satisfaction. In this case, individuals don't want paternalism, but do want institutions that produce benefits for the stakeholders (whether or not the paternalism is made explicit).

Open Question #9: What welfare criterion should we use to evaluate intertemporal tradeoffs?

A welfare criterion is needed in order to evaluate policy. While economists often rely on revealed preference, dynamically inconsistent preferences pose new challenges and how to evaluate welfare is disputed. Suppose an individual plans to contribute to a retirement savings plan next month (a patient choice), but when next month arrives changes her plans and spends the money (an impatient choice). Which choice makes the individual better off? Revealed preference has no obvious answer.

Normative analysis for intertemporal choice has been discussed through the lens of present bias (Laibson 1997). A common choice is to simply evaluate welfare from the long-run perspective (setting $\beta=1$), on the grounds that these are the preferences that are persistent; an alternative is the multi-self Pareto-criterion, which treats each self as different person (e.g., Laibson, Repetto, and Tobacman 1998).

There is a fundamental philosophical debate about the extent to which observed choices should be linked to normative preferences. Bernheim and Rangel (2009) and Bernheim (2016) argue that normative evaluation must be linked to choice and propose a welfare criterion that never overrules a choice; as a result, it is sometimes undiscerning (i.e., cannot rank two outcomes). In contrast, Beshears et al (2008) argue that *normative preferences* (that represent an individual's actual interests and that should be used for analysis) can diverge from revealed preferences (that

rationalize an action), and propose a variety of ways of identifying normative preferences. (See also Sugden 2004.) See the Handbook chapter in Volume 1 on Behavioral Public Economics (Bernheim and Taubinsky 2018) for more discussion of these issues.

The existence of paternalistic policies may reflect some important information about the appropriate welfare criteria: voters seem to demand and policy-makers seem to create paternalistic policies. It is an open question how informative the existence of paternalistic policies should be for informing welfare judgments. Policy can be mistaken, or policy may be insightful, or, most likely a bit of both. We might also be able to make inferences (in part) from the retirement savings systems and other human resource norms that have been developed by non-profit institutions that are concerned about the well-being of their employees (beyond the universal issues of recruitment and retention).

There has been a long tradition of not using revealed preference estimates of discount rates for normative analysis. Sen (1957) argues that the social discount rate should differ from the individual discount rate— “for, among other things, while the individuals die the nation does not.” Schelling (1995) notes the importance of conceptually distinguishing pure time preference for one’s own consumption from preferences for redistribution across generations, while Gollier (2013) highlights that the discount rate that should be used for an investment project depends on whether the project is financed by reducing current consumption or reallocating productive capital. Relatedly, Goulder and Williams (2012) similarly distinguish between a social-welfare discount rate and finance discount rate based on the marginal product of capital.

This normative discount rate has been particularly important for evaluating climate change and potential policy responses. Stern (2007, 2008) highlights how the social and private discount rates are distinct concepts; market interest rates reflect the private decisions of individuals but are not necessarily informative of the social discount rate, which also reflects ethical considerations.³⁹

3.10. Preference for Improving Sequences

Empirical Regularity #10: Individuals typically display both impatience and preferences for improvement over time.

Typically, individuals have (weakly) positive time preference: all else equal, they would rather have something earlier rather than later. However, there is one important exception: individuals seem to prefer sequences of outcomes that improve over time. For instance, Barsky et al (1997) show that participants prefer their real consumption path either to be flat or to rise over time (the modal subject prefers to it be flat, but the mean is a rising profile); when the real interest rate is zero, impatience and standard assumptions of separable discounted utility would predict a declining consumption profile. Similarly, Loewenstein and Sicherman (1991) show that

³⁹ However, even if the social discount rate is lower than the market interest rate, a policy maker might still wish to use the market interest rate as a return hurdle. Why invest in a public project with social return r when there exists a private project that can be used to generate a social return $r' > r$ that the government could pursue instead (Cropper and Laibson 1999)?

participants prefer an increasing wage profile,⁴⁰ and Chapman (2000) shows that subjects typically prefer improving sequences of health outcomes (e.g. headaches that get better over time rather than worse).

Open question #10: How do we integrate models of discounting with the other factors leading to a preference for improving sequences? When does a preference for an improving sequence become important relative to discounting?

Loewenstein and Prelec (1993) identified three related motivations for preferring an improving sequence: anticipatory utility, reference points, and a recency effect in evaluating experiences. Anticipatory utility can lead to preferences for improving sequences, if individuals prefer delaying pleasurable experiences so that they can savor them (Loewenstein 1987, see also Caplin and Leahy 2001); for instance, people prefer to delay receiving a kiss from a movie star by a few days, while also preferring to get through a painful shock immediately. Loss aversion with respect to reference points based on past consumption—adaptive reference points⁴¹— can also lead to a preference for increasing consumption, as can habit-based models (Ryder and Heal 1973, Abel 1990). Finally, because experiences may be evaluated with a “peak-end” rule (Kahneman et al 1993), individuals may prefer a better ending over a better beginning. It would be useful to quantitatively incorporate these motivations into a model of intertemporal choice that also involves discounting.

Moreover, Frederick and Loewenstein (2008) show that elicitation method matters for preferences over sequences: when choosing among sequences that differ in trend, people prefer increasing sequences; with other elicitation methods (allocation or pricing tasks), the preference for improving sequences disappears. If different motivations become salient with different tasks, which tasks best reflect economically relevant decisions?

4. Puzzles Without Associated Empirical Regularities

4.1. How Soon Is Now?

Open Question #11: How soon is “now”? How fast does value decline over time?

What is the duration of the “present” for present-focused models? Using a present-biased (i.e., quasi-hyperbolic) discounting model requires distinguishing between “now” and later. Of course, this discretization of time into the present versus the future is a simplification, but even a generalized hyperbolic discounting model requires a calibration of how fast value declines as you move forward.

Evidence on preferences over consumption suggests an extremely quick drop off in value overtime, on the order of minutes. McClure et al. (2007) examined choices over juice and water

⁴⁰ Read and Powell (2002) elicit verbal reasons for subjects’ choices over sequences over money, and find that subjects typically match income to desired consumption.

⁴¹ See DellaVigna et al (2017) and Thakral and Tô (2017) on adaptive reference points.

consumption for thirsty subject and found strong preference for drinking “now” versus more in 5 minutes, but substantially more patient choices between 20 and 25 minutes in the future.

Augenblick, Niederle, and Sprenger (2015) examine choices over when to exert effort and find present-bias for effort “now” (to be started in a few minutes) versus weeks later. Recently, Augenblick (2018) traces out discounting over effort tasks at short horizons (hours and days) and finds a discount factor of 0.94 for a delay of a few hours, 0.91 for one day, and 0.87 for one week. However, the interpretation of these results depends on the extent to which individuals can substitute effort across time; while today’s plans may be fixed, individuals may respond to increased requirements for effort in experiments in the future by reducing effort in other areas of their life.

A different picture emerges for discounting of monetary payments, which could be a result of a distinction between preferences over consumption versus financial flows. Balakrishnan, Haushofer, and Jakiela (2017) measure MEL choices using the convex time budget, comparing immediate payments, end of day payments, and future payments (using electronic payments for Kenyan subjects). They find little present-bias for end of day payments (about 4 hours delay) but do find present bias for immediate payments (β around 0.9).⁴² In contrast, structural models with annual periods (e.g. Angeletos et al. 2001) implicitly treat consumption anytime “this year” as immediate. The connection between the assumptions of these structural models and the experimental literature remains to be explored.

If there is limited present-bias beyond horizons of a few hours, the impact of present-bias on many decisions seems limited. For instance, purchasing goods online, even with next-day delivery, wouldn’t provide an immediate reward. Many expensive, durable goods (cars, computers) are only received at hours or days delay, and even clothing is not usually worn immediately after being purchased. However, there may be some transactional utility that makes the purchase of certain goods pleasurable in itself – i.e., buying a sports car or a special piece of clothing.

4.2. The Role of Temptation

Open Question #12: What types of decisions involve temptation? How quantitatively important is temptation (as opposed to present-bias)?

Intuition suggests that there is a cost to resisting temptation that is not captured in discounted utility models (whether they be time-consistent or present-biased), since in these discounted utility models the *outcome* (and not the cognitive process that gets the household to that endpoint) solely determines an individual’s welfare. Suppose you choose fruit as an afternoon snack; in a discounted utility model, your welfare is the same regardless of whether fruit was the only option, or you chose fruit over an unhealthy (but delicious) alternative. An alternative class of

⁴² Similarly, Andreoni and Sprenger (2012) find no present-bias for payments made via check at the end of the day. However, Augenblick, Niederle, and Sprenger (2015) made immediate monetary payments in cash and found β near 1 for money.

models allows for the possibility that resisting an unchosen alternative is costly (e.g. Gul and Pesendorfer (2001), Fudenberg and Levine (2006)).

Toussaert (2018) provides evidence on the existence of costly temptation. Participants in her experiment chose between either reading a story (the potentially “tempting” alternative) or completing a task for pay. About a third of her subjects preferred to limit their options for future work so that they could not read the story. This could be explained by either a desire to remove a tempting alternative that is costly to resist, or as a commitment device to change their future behavior. However, Toussaert shows that the desire to limit their options wasn’t driven by a desire to change their future behavior, since most participants predicted that they would do the task even if they had a choice. Moreover, consistent with a cost of resisting temptation, task performance was lower when participants were assigned to have the option to stop the task and read the story (even though they kept doing the task).

Other research also provides suggestive evidence that individuals want to limit their option set not merely to change their outcome but to reduce their costs of exercising self-control. Sadoff, Samek, & Sprenger (2015) examine the desire for commitment to healthy v. unhealthy foods in a food delivery service, and find that individuals who demand commitment are more likely to be dynamically consistent in their prior behavior. Similarly, examining gym attendance, Royer, Stehr, and Sydnor (2015) write, “we observe demand for commitment by those who appear on the surface not to need it.” This is hard to explain through the lens of present-biased discounted utility models but is naturally explained by models of costly temptation.

It is an open question how important temptation costs are for policy. Toussaert (2018) provides evidence of the existence of temptation, but the design had to be carefully constructed to rule out the alternative story in which participants demand commitment because they fear they might give in and choose the tempting alternative. Thus, her context has a tempting task where, perhaps unusually, people almost always resist temptation. Toussaert’s design also likely estimates a lower bound on the role of temptation, since it is joint test for both temptation preference *and* sophistication about those preferences: subjects have to predict that they will face costs of self-control. However, the literature examining time discounting and present-bias has typically found substantial naiveté (see Section 3.5).

It is difficult to distinguish costly temptation from other models such as present-bias (and note, the existence of costly temptation doesn’t rule out an important role for present-bias). Indeed, under certain conditions, Krusell, Kuruscu, and Smith (2010) show for strong enough costs of temptation (such that you always give in), behavior will be equivalent to that of present bias.⁴³ A smoking gun indicator for costly temptation is a willingness-to-pay to remove tempting things from your choice set that you wouldn’t choose anyway. But it is difficult to distinguish that motivation from willingness-to-pay for commitment if there is a small likelihood that you succumb to temptation. Future work may explore biometric data and subjective wellbeing measures to help measure pure temptation effects.

⁴³ Krusell, Kuruscu, and Smith (2010) show this equivalent for consumption choice under constant-relative-risk aversion preferences. It is unclear how general result can be made.

Large temptation costs could impact the evaluation of many policies, since in these models even options that are rarely if ever chosen could have substantial welfare impacts. Hence, smoking bans might increase the welfare of former smokers who have quit, but who no longer have to resist the temptation of smoking. The availability of social media on smart phones might reduce welfare, even when people aren't using their cell phones. Finally, deadlines might be valuable because they lower the self-control costs of working on a project, even if they don't ultimately change when the project is completed.

4.3. Other mechanisms and complementary psychological conceptions

Open Question #13: How important are a variety of mechanisms for intertemporal choice, including probability weighting, trust, and heuristics, both in the lab and in the field? What role do alternative psychological conceptions play in intertemporal choice?

Discounted utility (as well as other models of present-focus) is fundamentally a theory of time preference. Yet intertemporal choices also reflect other factors, such as risk, heuristics, or even framing that are distinct from pure time preference. We do not yet have a clear picture of how quantitatively important these other factors are.

Risk

An alternative motivation for discounting the receipt of future rewards is that the reward may disappear—e.g. the payment may not be made, or the individual may die, etc.⁴⁴ A common concern in interpreting MEL tasks is whether subjects trust that payments will be delivered, and experimenters attempt to maintain high reliability in when payments are received. Risk that a future reward might disappear can lead to discounting apart from pure time preference, as discussed in Section 1.4.

Other work moves beyond expected utility models. Halevy (2008) argues that the present is different from the future in that the future is uncertain; in non-expected utility models, this can lead to diminishing impatience (see also Saito 2011). Along these lines, Epper, Fehr-Duda, and Bruhin (2011) show experimental evidence that individuals who are prone to more probability distortions also discount money more hyperbolically, and show that, since future events are uncertain, probability distortions can lead to hyperbolic discounting.

Other work shows interactions between risk and time that are not predicted by expected utility. For instance, Anderson and Stafford (2009) find that the presence of risk makes individuals less patient (adjusting for the level of risk); see also Ahlbrecht and Weber (1997). Andreoni and Sprenger (2012b) show that participants seem to display a preference for certainty, though Cheung (2015) shows that those results are at least partially driven by a motive for diversifying across dates and thereby hedging payment risk (because the risks are independent); see also Miao and Zhong (2015) and a reply by Andreoni and Sprenger (2015).

⁴⁴ McGuire and Kable (2013) argue that behavior in the original marshmallow task are driven by trust and changing beliefs about how long it would be necessary to wait for the larger reward.

Heuristics

It may be cognitively difficult to discount. For instance, evidence on exponential growth bias (Stango and Zinman 2009) indicates that people get exponentiation wrong and tend to linearize exponential functions; this poses a challenge for discounted utility models. Given cognitive difficulties, individuals may rely on heuristics. Various models of intertemporal choice heuristics have been proposed in which choice is a function of proportional and absolute differences on time and reward dimension, as well as other factors (Read et al. 2013, Ericson et al 2015). Ericson et al. (2015) conduct a cross-validated intertemporal choice model comparison showing that heuristic models predict better out of sample than discounted utility models, though Wulff and van den Bos (2017) show that some of their results depend on the model of noise and individual heterogeneity.

Heuristic models of intertemporal choice seem promising but are currently limited in their applicability. It remains an open question whether heuristic models generalize beyond lab-based MEL choices. Can economists use them to explain any other behaviors? Moreover, it may be fruitful to develop a model that nests both heuristics and discounted utility, allowing researchers to estimate a parameter capturing the extent of deviation from expected utility.

Other Theories

Alongside the economics literature is a psychological literature examining intertemporal choice, and it may be fruitful to examine the relationship between the psychological and economic models.⁴⁵ For instance, personality psychology characterizes individuals not via their discount factor, but via concepts such as conscientiousness and grit.⁴⁶ Conscientiousness is one of the “big five” personality measures, and may include some aspects of self-control (Roberts et al 2005). Grit, defined as perseverance and passion for long-term goals, has been argued to be distinct from self-control (Duckworth and Gross 2014) and associated with measures of life success (Duckworth et al 2007). How does grit translate into an economic model? Can grit be captured by a utility function that places less weight on the cost of effort? Or is grit an increased ability to forecast and make salient future benefits of present costly actions?

Other work examines factors that affect the construction of intertemporal choice. Query theory (Weber et al. 2007, Appelt, Hardisty, and Weber 2011) can be used to understand the process of making an intertemporal choice. Weber et al. (2007) show that manipulating the order in which individuals think of different reasons for their choice can lead to more or less impulsivity. Another line of work examines the role of emotion in intertemporal choice, showing that sadness increases impatience (Lerner, Li, and Weber 2013) and that gratitude decreases impatience (DeSteno, Li, Dickens, and Lerner 2014).

⁴⁵ For a review of related psychological research on behavior change, see Duckworth, Laibson, and Milkman (forthcoming).

⁴⁶ See also discussion of the connection between economics and personality psychology in Borghans, Duckworth, Heckman and ter Weel's (2008), and Becker et al. (2012).

4.4. Stability and Domain Generality

Open Question #14: How stable are time preferences? How general are they across domains?

In discounted utility models, time preference applies to utility, and so the same discount factor applies to all domains of choice. Yet it is an open question whether different types of consumption are discounted in similar ways.

Cohen, Ericson, Laibson, and White (2016) review how predictive money-earlier-or-later (MEL) studies are for other domains. Many studies show that impatience in MEL tasks is weakly correlated with other behaviors. For instance, Meier and Sprenger (2011) find that MEL discounting is correlated with individuals' credit scores, and a series of papers shows that children's MEL choices predict future behavior (Bettinger and Slonim 2007, Castillo et al. 2011, Sutter et al. 2013, Golsteyn, Gronqvist and Lindahl 2013).⁴⁷ Aggregating field behaviors into single indices increases the correlation with MEL behavior (Chabris et al 2008). However, as discussed previously, preferences over the timing of money receipt can (and should in an optimizing model) diverge from preferences over other rewards, since receipt of money is not the same as consumption.

There is limited work comparing consumption-based measures of time preference to field behavior. A classic non-monetary task is the marshmallow test (Mischel and Ebbesen 1970, Mischel, Shoda, and Rodriguez 1989), which is correlated with personality scales that measure self-regulatory capacity; it has some predictive power for early-life outcomes (Shoda, Mischel, and Peake 1990, Watts, Duncan, and Quan 2018) and no predictive power for mid-life outcomes (Benjamin et al 2018). Burks et al. (2012) develop an analogue of the marshmallow task, in which individuals trade off waiting longer for monetary compensation. They conclude that while this measure is correlated with smoking, it “does not predict outcomes very well”, and that a quasi-hyperbolic model measured from MEL choices does a better job.

It would be interesting to examine how time preferences over different domains of consumption are correlated, but evidence is limited here. A number of studies examine money and another domain (e.g. health and money in Chapman 1996 or effort and money in Augenblick, Niederle, and Sprenger 2015). Other studies examine field behaviors that could be a function of time preferences but also other factors (e.g. smoking, obesity; see Chabris et al 2008, Falk et al 2015, 2016). The best evidence on discounting of different types of goods comes from developing countries. Ubfal (2016) estimates good specific discount rates (e.g. for sugar, rice, meat) and rejects hypothesis that discount rates are the same across goods.

Correlations between measured time preference (whatever the method) and field behaviors have multiple interpretations, as behaviors in both the lab and field are functions not only of time preference, but also cognitive ability, knowledge, beliefs, and preferences. For instance, more

⁴⁷ Another type of correlation is between measures of restlessness, aggression, or impulsivity in children and their future outcomes (Cadena and Keys 2012, Moffit et al 2011); it is unclear how to interpret these measures in a model of discounting utility.

patient MEL behavior is correlated with measures of cognitive ability (Burks et al. 2009, Dohmen et al. 2010, and Benjamin, Brown, and Shapiro 2013). Thus, unconditional correlations between MEL behavior and other behaviors include some effect of cognitive ability. Why cognitive ability and MEL behavior are correlated is another puzzle. One explanation is that more patience leads to more investment in human capital. An alternative explanation is that higher cognitive ability makes it more likely that individuals treat MEL decisions “rationally” as a financing decision rather than a consumption decision; their choices thus reveal their market interest rate rather than a discount rate for utils.

Finally, the role of noise in eliciting or constructing preference may be important. Test-retest correlations of discount rates measured from MEL behavior are similar to other personality measures: 0.6 to 0.8 in Kirby (2009), and 0.5 in Meier and Sprenger (2015). Yet it is also rare that a discounting model strongly predicts an individual’s choices; the form of noise used to estimate discounting parameters can be quite important. For instance, it may matter whether noise is additive (so that as rewards are scaled up, an individual is less likely to make a “mistake”) or multiplicative (so that mistakes are equally likely for high stakes choices), or some hybrid case.

4.5. Malleability and self-management

Open Question #15: How malleable is time preference? How effective are self-management techniques?

Commitment is one self-management strategy. Yet there are other ways individuals might alter their intertemporal choices.

Early childhood interventions have been shown to have large impacts on later life outcomes; much of this effect comes from non-cognitive skills, potentially including self-regulation and/or time preference (Heckman, Stixrud, and Urzua 2006). However, we are unaware of direct evidence showing the effect of early childhood interventions on conventional measures of time preference (MEL or real consumption choices). Heckman, Pinto, and Savelyev (2013) show that an early childhood intervention improved academic motivation (e.g. “showing initiative”) and externalizing behaviors (e.g. disruption of class, aggressive behaviors) that seem plausibly related to time preference or self-control.

One hypothesis is that cognitive load can affect impatience, a prediction that can be produced by cost-of-self-control models. The evidence here is mixed. Deck and Jahedi (2015) provide experimental evidence that increasing cognitive load leads to more impatience over money, but not over consumption goods (see also Shiv and Fedorikhin 1999, Hinson, Jameson, and Whitney 2003).⁴⁸ Closely connected is the hypothesis that scarcity leads to myopic or impatient behavior, supported by experimental studies (Shah, Mullainathan, and Shafir 2012; see also Haushofer and Fehr 2014). Carvalho, Meier, and Wang (2015) examine choices of low-income individuals

⁴⁸ Exogenously increasing time for deliberation has been shown to lead to more patient decisions (Dai and Fishbach 2013, Imas, Kuhn, and Mironova 2017).

before versus after payday (generating a short-lived change in financial pressure). They show that before payday, choices are more impatient over money, but are *not* more impatient over real effort tasks; this is naturally explained via liquidity constraints, rather than changing time preferences.

Other work has shown that a “growth mindset” (a belief that intelligence is malleable rather than fixed) predicts greater academic performance (Dweck 1986, Blackwell et al., 2007). Meta-analyses have found small effects of mindset interventions (Sisk et al 2018) with larger effects for low-socio-economic status students. A large-scale replication of a growth-mindset intervention found significant, though small effects (Yeager et al. 2018). While the growth mindset is likely best captured as a change in beliefs about self-efficacy and the connection between effort and reward, it might also impact decisions about self-control. (See also our discussion of “grit” in Section 13.)

4.6. Retirement saving adequacy

Open question #16: Are households saving optimally for retirement?

Since the work of Hamermesh (1984) and Mariger (1987) there has been an active literature studying consumption dynamics around retirement and the closely related question of the adequacy of household retirement savings.⁴⁹ There are strong advocates of two positions: that households save too little and that households save optimally. There is also a small set of papers that argue that many households are tightwads and accordingly save too much (e.g., Rick, Cryder, and Loewenstein, 2007).

Many studies have shown that, on average, consumption expenditure falls around retirement (e.g., Bernheim et al 2001; Angeletos et al 2001; Haider and Stephens Jr., 2007; Olafsson and Pagel 2018). Many rational mechanisms might explain this drop, including a decline in work-related expenditures, a shift toward home production (i.e., leisure-consumption substitution), and a greater propensity to shop for low-cost consumption goods (e.g., more time for coupon clipping or travelling a greater distance to purchase goods on sale).

Bernheim et al. (2001) show that the decline in consumption expenditure at retirement is not associated with the decline in work-related expenditures. Relatedly, Ganong and Noel (2017) show that work-related expenditures do not account for a substantial change in expenditure when a member of a household separates from an employer.

A hypothesized shift toward home production at retirement, which would smooth *total* market and non-market consumption even as market-based consumption *expenditure* fell, was first supported by the finding that *caloric* consumption did not drop at retirement (Aguiar and Hurst 2005). However, subsequent work has challenged that finding with new data (Stephens and Toohey (2018)). It now appears that caloric consumption does fall on average at retirement, though it is still not clear how to interpret that finding. It is possible that retirees expend less

⁴⁹ For a review of the literature on consumption and savings, see Attanasio and Weber (2010).

energy than employed people (holding age fixed). It would be useful to longitudinally follow the caloric intake of a person who is not in the labor force and passes from a period in which their spouse is employed to a period in which their spouse is retired. Of course, even this empirical strategy presents confounds if the retirement of one's spouse changes one's own level of physical activity.

There is evidence for the hypothesized mechanism that retirees purchase less expensive goods in retirement (holding quality fixed): e.g., Aguiar and Hurst (2007), Hurd and Rohwedder (2013) and Agarwal et al (2015). But it is not clear whether this behavior reflects a rational plan for lifecycle behavior (i.e., another example of rational substitution between leisure and consumption expenditure) or the consequence of financial pressure/distress that leads retirees to search for bargains.

There is also a literature that studies wealth accumulation. For example, Scholz et al (2006) use a structural model to analyze the optimality of lifecycle wealth formation and report that the vast majority of households are saving optimally. Here too there are challenges of interpretation. Key parameters make a large difference in these models (Skinner 2007). For example, if households experience large incremental costs when they have children, then it is optimal for expenditure to fall sharply when children leave the home, and by extension for balance sheet wealth for households on the brink of retirement (who are largely empty nesters) to be low. For evidence that consumption expenditure changes relatively little when children start to live independently, see Dushi et al. (2015).

In surveys, households overwhelmingly report that they are saving less than they think they should (e.g., Bernheim 1991; Bernheim 1995; Farkas and Johnson 1997; Choi et al. 2002), though it is hard to interpret this attitudinal measure. Such perceived undersaving is consistent with models of present-bias, which have been used to explain wealth accumulation, credit card borrowing, and consumption-income comovement, including the drop in consumption expenditure at retirement (Angeletos et al 2001; Laibson, Repetto, and Tobacman 2003; Laibson, Maxted, Repetto, and Tobacman 2017).

The extreme sensitivity between savings behavior and institutional nudges (e.g., auto-enrollment) suggests that savings behavior is not consistent with a rational actor model (e.g., Madrian and Shea 2001). But this does not necessarily imply that households are undersaving on average. It is possible that paternalistic policies (like Social Security and other pensions) are structured so that on average households accumulate wealth optimally, even though each individual household is not necessarily at its own private optimum (Fadlon and Laibson, 2018).

5. Conclusion

A set of key stylized facts have been documented in the literature on intertemporal choice. These facts are largely at odds with the once-dominant exponential discounted utility model. A variety of models with the feature of present-focus have been developed to describe how humans make intertemporal choices. While in many cases these models make similar predictions, they have

important distinctions that can affect the advice we give individuals, the policies we develop, and the welfare judgments we make.

As a result, despite a robust literature, there is a need for new empirical evidence to address the important open questions that we identify and to distinguish among various present-focused models. There is thus room for much new work in the field of intertemporal choice, which may be furthered by new experimental paradigms. Much of the laboratory research we review relies on monetary tradeoffs, but the connection between these money choices and experienced utility is itself unclear and a subject of ongoing research. Real effort laboratory tasks are a promising methodology that poses considerable advantages over experiments based on time-linked monetary payoffs (e.g. Augenblick, Niederle and Sprenger 2015, Augenblick and Rabin 2018, Fedyk 2018), but these methods are also imperfect. Other sources of data, such as household balance sheet data (Laibson, Maxted, Repetto, and Tobacman, 2017), spending transactions (e.g. Ganong and Noel 2017; Kuchler and Pagel 2018) or caloric intake (Aguiar and Hurst 2005; Stephens and Toohey 2018) will complement the methods that have been deployed in the historical literature. Ultimately, there probably isn't a right answer to our methodological quest. Progress will likely come from many different methods that jointly paint a more complete and compelling picture of our intertemporal preferences.

6. References

- Abadie, A., Gay, S. (2006). "The impact of presumed consent legislation on cadaveric organ donation: A cross-country study". *Journal of Health Economics* 25 (4), 599-620.
- Abdellaoui, M., Attema, A.E., Bleichrodt, H. (2010). "Intertemporal Tradeoffs for Gains and Losses: An Experimental Measurement of Discounted Utility". *The Economic Journal* 120 (545), 845–66.
- Abdellaoui, M., Bleichrodt, H., l'Haridon, O., Paraschiv, C. (2013). "Is there one unifying concept of utility? An experimental comparison of utility under risk and utility over time". *Management Science* 59 (9), 2153-2169.
- Abel, A.B. (1990). "Asset Prices under Habit Formation and Catching up with the Joneses." *The American Economic Review* 80 (2), 38–42.
- Acland, D., Levy, M.R. (2015). "Naiveté, Projection Bias, and Habit Formation in Gym Attendance". *Management Science* 61 (1), 146-160.
- Agarwal, S., Pan, J., Qian, W. (2015). "The Composition Effect of Consumption around Retirement: Evidence from Singapore". *American Economic Review* 105 (5), 426-431.
- Aguiar, M., Hurst, E. (2005). "Consumption versus Expenditure". *Journal of Political Economy* 113 (5), 919-948.
- Aguiar, M., Hurst, E. (2007). "Life-Cycle Prices and Production". *American Economic Review* 97 (5), 1533-1559.
- Ahlbrecht, M., Weber, M. (1997). "An Empirical Study on Intertemporal Decision Making Under Risk." *Management Science* 43 (6), 813–26.
- Ahn, D.S., Iijima, R., Sarver, T. (2017). *Naivite About Temptation and Self-Control: Foundations for Naive Quasi-Hyperbolic Discounting*. Cowles Foundation Discussion Papers 2099.
- Ainslie, G. (1975). "Specious reward: a behavioral theory of impulsiveness and impulse control". *Psychological Bulletin* 82 (4), 463-496.
- Aizer, A. (2007). "Public health insurance, program take-up, and child health". *The Review of Economics and Statistics* 89 (3), 400-415.
- Akerlof, G.A. (1991). "Procrastination and Obedience". *American Economic Review* 81 (2), 1-19.

- Albrecht, K., Volz, K.G., Sutter, M., Laibson, D.I., von Cramon, D.Y. (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.
- Allcott, H., Greenstone, M. (2012). "Is There an Energy Efficiency Gap?". *Journal of Economic Perspectives* 26 (1), 3-28.
- Allcott, H., Wozny, N. (2014). "Gasoline prices, fuel economy, and the energy paradox". *Review of Economics and Statistics* 96 (5), 779-795.
- Alsan, M., Beshears, J., Armstrong, W., Choi, J.J., Madrian, B.C., Nguyen, M.L.T., ... & Marconi, V.C. (2017). "A commitment contract to achieve virologic suppression in poorly adherent patients with HIV/AIDS." *Aids* 31 (12), 1765-1769.
- Amador, M., Werning, I., & Angeletos, G. M. (2006). "Commitment vs. flexibility." *Econometrica*, 74(2), 365-396.
- Andersen, S., Harrison, G.W., Lau, M.I., Rutström, E.E. (2008). "Eliciting Risk and Time Preferences". *Econometrica* 76 (3), 583–618.
- Anderson, L.R., Stafford, S.L. (2009). "Individual Decision-Making Experiments with Risk and Intertemporal Choice". *Journal of Risk and Uncertainty* 38 (1), 51–72.
- Andreoni, J., Aydin, D., Barton, B., Bernheim, B.D., Naecker, J. (2016). When Fair Isn't Fair: Understanding Choice Reversals Involving Social Preferences. National Bureau of Economic Research Working Paper No. 25257.
- Andreoni, J., Barton, B., Bernheim, B.D., Aydin, D., Naecker, J. (2016). "When Fair Isn't Fair: Sophisticated Time Inconsistency in Social Preferences". Working Paper.
- Andreoni, J., Gravert, C., Kuhn, M.A., Saccardo, S., Yang, Y. (2018). Arbitrage or Narrow Bracketing? On Using Money to Measure Intertemporal Preferences. National Bureau of Economic Research Working Paper No. 25232.
- Andreoni, J., Kuhn, M.A., Sprenger, C. (2015). "Measuring Time Preferences: A Comparison of Experimental Methods". *Journal of Economic Behavior and Organization* 116, 451-464.
- Andreoni, J., Sprenger, C. (2012a). "Estimating Time Preferences from Convex Budgets". *American Economic Review* 102 (7), 3333–3356.
- Andreoni, J., Sprenger, C. (2012b). "Risk Preferences Are Not Time Preferences". *The American Economic Review* 102 (7), 3357–3376.
- Andreoni, J., Sprenger, C. (2015). "Risk Preferences Are Not Time Preferences: Reply". *The American Economic Review* 105 (7), 2287–2293.
- Angeletos, G., Laibson, D., Repetto, A., Tobacman, J., Weinberg, S. (2001). "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation". *The Journal of Economic Perspectives* 15 (3), 47–68.
- Appelt, K.C., Hardisty, D.J., Weber, E.U. (2011). "Asymmetric Discounting of Gains and Losses: A Query Theory Account". *Journal of Risk and Uncertainty* 43 (2), 107–126.
- Ariely, D., Wertenbroch, K. (2002). "Procrastination, Deadlines, and Performance: Self-Control by Precommitment". *Psychological Science* 13 (3), 219–224.
- Aristotle. (2000). *Nicomachean Ethics*, 2nd edition. Translated, with introduction, by Terence Irwin. Hackett Publishing, Indianapolis/Cambridge.
- Ashraf, N., Camerer, C.F., Loewenstein, G. (2005). "Adam Smith, Behavioral Economist". *The Journal of Economic Perspectives* 19 (3), 131–145.
- Ashraf, N., Karlan, D., Yin, W. (2006). "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines". *The Quarterly Journal of Economics* 121 (2), 635–72.
- Attanasio, O.P., Weber, G. (2010). "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy." *Journal of Economic Literature* 48: 693-751.
- Augenblick, N. (2018). Short-Term Time Discounting of Unpleasant Tasks. Working Paper.
- Augenblick, N., Niederle, M., Sprenger, C. (2015). "Working Over Time: Dynamic Inconsistency in Real Effort Tasks". *Quarterly Journal of Economics*. 30 (3), 1067-1115.

- Augenblick, N., Rabin, M. (2018). An Experiment on Time Preference and Misprediction in Unpleasant Tasks. *Review of Economic Studies*, forthcoming.
- Azfar, O. (1999). "Rationalizing Hyperbolic Discounting". *Journal of Economic Behavior & Organization* 38 (2), 245-252.
- 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.
- Balakrishnan, U., Haushofer, J., Jakiela, P. (2017). How soon is now? Evidence of present bias from convex time budget experiments. National Bureau of Economic Research Working Paper No. 23558.
- Banerjee, A.V., Duflo, E., Glennerster, R. (2010). "Improving immunisation coverage in rural India: clustered randomised controlled evaluation of immunisation campaigns with and without incentives". *BMJ* 340.
- Banerjee, A., Mullainathan, S. (2010). The shape of temptation: Implications for the economic lives of the poor. National Bureau of Economic Research Working Paper No. 15973.
- Barro, R.J. (1999). "Ramsey meets Laibson in the neoclassical growth model". *The Quarterly Journal of Economics* 114 (4), 1125-1152.
- Barsky, R.B., Juster, F.T., Kimball, M.S., Shapiro, M.D. (1997). "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study". *Quarterly Journal of Economics* 112 (2), 537-579.
- Baucells, M., Heukamp, F.H. (2012). "Probability and Time Trade-Off". *Management Science* 58 (4), 831-842.
- Baumeister, R.F., Bratslavsky, E., Muraven, M., Tice, D.M. (1998). "Ego Depletion: Is the Active Self a Limited Resource?". *Journal of Personality and Social Psychology* 74 (5), 1252-1265.
- Becker, G.M., Degroot, M.H., Marschak, J. (1964). "Measuring Utility by a Single-Response Sequential Method". *Behavioral Science* 9 (3), 226–32. doi:10.1002/bs.3830090304.
- Becker, G.S., Murphy, K.M. (1988). "A theory of rational addiction". *Journal of Political Economy* 96 (4), 675-700.
- Benartzi, S. (2001). "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock". *Journal of Finance* 56 (5), 1747-1764.
- Benartzi, S., Previtro, A., Thaler, R.H. (2011). "Annuitization puzzles". *Journal of Economic Perspectives* 25 (4), 143-164.
- Benhabib, J., Bisin, A., Schotter, A. (2010) "Present-bias, quasi-hyperbolic discounting, and fixed costs". *Games and Economic Behavior* 69 (2), 205-223.
- Benjamin, D.J., Brown, S.A., Shapiro, J.M. (2013). "Who is 'Behavioral'? Cognitive Ability and Anomalous Preferences". *Journal of the European Economic Association* 11 (6), 1231-1255. doi:10.1111/jeea.12055.
- Benjamin, D.J., Laibson, D., Mischel, W., Peake, P.K., Shoda, Y., Steiny, A., Wilson, N.L. (2018). Predicting Mid-Life Capital Formation with Life-Course Measures of Self-Regulation. Working Paper.
- Bernheim, B.D. (1995). "Do Household Appreciate Their Financial Vulnerabilities? An Analysis of Actions, Perceptions, and Public Policy". In: *Tax Policy and Economic Growth*. American Council for Capital Formation, Washington, DC.
- Bernheim, B.D. (1991). *The vanishing nest egg: Reflections on saving in America*. Priority Press Publications, New York.
- Bernheim, B.D. (2016). "The Good, the Bad, and the Ugly: A Unified Approach to Behavioral Welfare Economics". *Journal of Benefit-Cost Analysis* 7 (1), 12–68. doi:10.1017/bca.2016.5.
- Bernheim, B.D., Fradkin, A., Popov, I. (2015). "The Welfare Economics of Default Options in 401(k) Plans". *American Economic Review* 105 (9), 2798-2837.
- Bernheim, B.D., Meer, J., Novarro, N.K. (2016). "Do Consumers Exploit Commitment Opportunities? Evidence from Natural Experiments Involving Liquor Consumption". *American Economic Journal: Economic Policy* 8 (4), 41-69.
- Bernheim, B.D., Rangel, A. (2004). "Addiction and cue-triggered decision processes". *American Economic Review* 94 (5), 1558-1590.

- Bernheim, B.D., Rangel, A. (2009). "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics". *Quarterly Journal of Economics* 124 (1), 51–104. doi:10.1162/qjec.2009.124.1.51.
- Bernheim, B.D., Ray, D., Yeltekin, S. (2015). "Poverty and Self-Control". *Econometrica* 83 (5), 1877–1911. doi:10.3982/ECTA11374.
- Bernheim, B.D., Skinner, J., Weinberg, S. (2001). "What Accounts for the Variation in Retirement Wealth among U.S. Households?". *American Economic Review* 91 (4), 832–857. doi:10.1257/aer.91.4.832.
- Bernheim, B. D., & Taubinsky, D. (2018). "Behavioral public economics." In: Bernheim, B.D., Laibson D., DellaVigna S. (Eds.), *Handbook of Behavioral Economics*. North-Holland, Amsterdam.
- Beshears, J., Choi, J.J., Harris C., Laibson, D., Madrian, B.C. (2018). "Optimal Illiquidity." Mimeo.
- Beshears, J., Choi, J.J., Hurwitz, J., Laibson, D., Madrian, B.C. (2015). "Liquidity in Retirement Savings Systems: An International Comparison". *American Economic Review* 105 (5): 420-425.
- Beshears, J., Choi J.J., Laibson D., Madrian B.C. (2008). "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States". In: Kay, S.J., Sinha T. (Eds.), *Lessons from Pension Reform in the Americas*. Oxford University Press, Oxford.
- Beshears, J., Choi J.J., Laibson D., Madrian B.C. (2013). "Simplification and saving". *Journal of Economic Behavior & Organization* 95, 130-145.
- Beshears, J., Choi J.J., Laibson D., Madrian B.C. (2018). "Behavioral Household Finance". In: Bernheim, B.D., Laibson D., DellaVigna S. (Eds.), *Handbook of Behavioral Economics*. North-Holland, Amsterdam.
- Beshears, J., Choi J.J., Laibson D., Madrian B.C., Sakong, J. (2018). "Self Control and Liquidity: How to Design a Commitment Contract." Working Paper.
- Beshears, J., Choi J.J., Laibson D., Madrian B.C., Zeldes, S.P. (2014). "What makes Annuitization More Appealing?". *Journal of Public Economics* 116, 2-16.
- Bettinger, E.P., Long, B.T., Oreopoulos, O., Sanbonmatsu, L. (2012). "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment". *The Quarterly Journal of Economics* 127 (3), 1205-1242.
- Bettinger, E., Slonim, R. (2007). "Patience among Children." *Journal of Public Economics* 91 (1–2), 343–63. doi:10.1016/j.jpubeco.2006.05.010.
- Bhargava, S., Manoli, D. (2015). "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11), 3489-3529.
- Bisin, A., Hyndman, K. (2018). Present-Bias, Procrastination and Deadlines in a Field Experiment. National Bureau of Economic Research Working Paper No. 19874.
- Blackwell, L., Trzesniewski, K., Dweck, C. S. (2007). "Implicit theories of intelligence predict achievement across an adolescent transition: A longitudinal study and an intervention". *Child Development* 78, 246–263.
- Blumenstock, J., Callen, M., Ghani, T. (2018). "Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan". *American Economic Review* 108 (10), 2868-2901.
- Böhm-Bawerk, Eugen, von. (1890). *Capital and Interest: A Critical History of Economical Theory*. Macmillan, London.
- Borghans, L., Duckworth, A.L., Heckman, J.J., ter Weel, B. (2008). "The Economics and Psychology of Personality Traits". *Journal of Human Resources* 43, 972-1059.
- Brocas, I., Carrillo, J.D. (2008a). "Theories of the Mind". *The American Economic Review* 98 (2), 175–80.
- Brocas, I., Carrillo, J.D. (2008b). "The Brain as a Hierarchical Organization." *The American Economic Review* 98 (4), 1312–46. doi:10.1257/aer.98.4.1312.
- Brocas, I., Carrillo, J.D. (2012). "From Perception to Action: An Economic Model of Brain Processes." *Games and Economic Behavior* 75 (1), 81–103. doi:10.1016/j.geb.2011.10.001.
- Broda, C., Parker, J.A. (2014). The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption. National Bureau of Economic Research Working Paper No. 20122.
- Brown, D.J., Lewis, L.M. (1981). "Myopic Economic Agents". *Econometrica* 49 (2), 359-68.

- Brown, J.R., Kapteyn, A., Luttmer, E.F., Mitchell, O.S. (2017). "Cognitive constraints on valuing annuities". *Journal of the European Economic Association* 15 (2), 429-462.
- Brown, J.R., Liang, N., Weisbenner S. (2007). "Individual Account Investment Options and Portfolio Choice: Behavioral Lessons from 401(k) Plans". *Journal of Public Economics* 91 (10), 1992-2013.
- Brown, J.R., Previtro, A. (2016). Saving for retirement, Annuities and Procrastination. Working Paper.
- Bryan, G., Karlan, D., Nelson, S. (2010). "Commitment Devices." *Annual Review of Economics* 2 (1), 671-98. doi:10.1146/annurev.economics.102308.124324.
- Burger, N., Charness, G., Lynham, J. (2011). "Field and Online Experiments on Self-Control." *Journal of Economic Behavior and Organization* 77 (3), 393-404. doi:10.1016/j.jebo.2010.11.010.
- Burks, S.V., Carpenter, J.P., Goette, L., Rustichini, A. (2009). "Cognitive Skills Affect Economic Preferences, Strategic Behavior, and Job Attachment". *Proceedings of the National Academy of Sciences* 106 (19), 7745-7750.
- Burks, S., Carpenter, J., Götte, L., Rustichini, A. (2012). "Which Measures of Time Preference Best Predict Outcomes: Evidence from a Large-Scale Field Experiment." *Journal of Economic Behavior & Organization* 84 (1), 308-20. doi:10.1016/j.jebo.2012.03.012.
- Busse, M.R., Knittel, C.R., Zettelmeyer, F. (2013). "Are Consumers Myopic? Evidence from New and Used Car Purchases." *The American Economic Review* 103 (1), 220-256. doi:10.1257/aer.103.1.220.
- Cadena, B.C., Keys, B.J. (2012). Human Capital and the Lifetime Costs of Impatience. SSRN Scholarly Paper ID 1674068, Social Science Research Network.
- Camerer, C.F., Cohen, J.D., Fehr, E., Glimcher, P.W., Laibson, D. (2015). "Neuroeconomics." In: Kagel, J.H., Roth, A.E. (Eds.), *Handbook of Experimental Economics*, vol 2. Princeton University Press, Princeton and Oxford.
- Campbell, J.Y., Cochrane, J.H. (1999). "By force of habit: A consumption-based explanation of aggregate stock market behavior". *Journal of Political Economy* 107 (2), 205-251.
- Caplin, A., Leahy, J. (2001). "Psychological expected utility theory and anticipatory feelings". *The Quarterly Journal of Economics* 116 (1), 55-79.
- Carroll, G.D., Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A. (2009) "Optimal defaults and active decisions". *The Quarterly Journal of Economics* 124 (4), 1639-1674.
- Carvalho, L., Meier, S., Wang, S. (2016). "Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday". *American Economic Review* 106 (2), 260-284.
- Casaburi, L., and Macchiavello R. (forthcoming) "Demand and Supply of Infrequent Payments as a Commitment Device: Evidence from Kenya", *American Economic Review*.
- Castillo, M., Ferraro, P., Jordan, J.L., Petrie, R. (2011). "The Today and Tomorrow of Kids: Time Preferences and Educational Outcomes of Children." *Journal of Public Economics* 95 (11), 1377-85.
- Cassidy, R. (2017). "Are the Poor Really so Present-Biased? Evidence from a Field Experiment in Rural Pakistan." IFS Working Paper WP18/24.
- Chabris, C.F., Laibson, D., Morris, C.L., Schuldt, J.P., Taubinsky, D. (2008). "Individual Laboratory-Measured Discount Rates Predict Field Behavior". *Journal of Risk and Uncertainty* 37 (2-3), 237-269.
- Chabris, C.F., Laibson, D.I., Schuldt, J.P. (2010). "Intertemporal choice". In: Durlauf, S.N., Blume, L.E. (Eds.), *Behavioural and Experimental Economics*. Palgrave Macmillan, London.
- Chakraborty, A. (2017). "Present Bias". Working Paper.
- Chapman, G.B. (1996). "Temporal Discounting and Utility for Health and Money". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22 (3), 771-791.
- Chapman, G.B. (2000). "Preferences for improving and declining sequences of health outcomes". *Journal of Behavioral Decision Making* 13 (2), 203-218.
- Cheung, S. (2015). "Risk Preferences are Not Time Preferences: On the Elicitation of Time Preference under Conditions of Risk: Comment." *American Economic Review* 105 (7), 2242-2260.
- Cheung, S. L. (2016). "Recent developments in the experimental elicitation of time preference". *Journal of Behavioral and Experimental Finance* 11, 1-8.

- Choi, J.J., Laibson, D., Madrian, B.C. (2009). "Mental Accounting in Portfolio Choice: Evidence from a Flypaper Effect". *American Economic Review* 99 (5), 2085-2095.
- Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A. (2002). "Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance". In: Poterba, J.M. (Ed.), *Tax Policy and the Economy*, vol. 16, National Bureau of Economic Research. MIT Press, Cambridge.
- Choukhmane, T. (2018). *Default Options and Retirement Savings Dynamics*. Working Paper.
- Chung, S., Herrnstein, R.J. (1967). "Choice and Delay of Reinforcement." *Journal of the Experimental Analysis of Behavior* 10 (1), 67–74. doi:10.1901/jeab.1967.10-67.
- Cohen, J.D., Ericson, K.M., Laibson, D., White, J.M. (2016). *Measuring Time Preferences*. National Bureau of Economic Research Working Paper No. 22455.
- Coile, C., Diamond, P., Gruber, J., Jousten, A. (2002). "Delays in Claiming Social Security Benefits." *Journal of Public Economics* 84 (3), 357–385. doi:10.1016/S0047-2727(01)00129-3.
- Coller, M., Williams, M.B. (1999). "Eliciting Individual Discount Rates." *Experimental Economics* 2 (2), 107–27. doi:10.1023/A:1009986005690.
- Commons, M.L., Woodford, M., Ducheny, J.R. (1982). "How Reinforcers Are Aggregated in Reinforcement-Density Discrimination and Preference Experiments". In: Commons, M.L., Herrnstein, R.J., Rachlin, H. (Eds.), *Quantitative Analyses of Behavior: Vol. 2, Matching and Maximizing Accounts*. Ballinger, Cambridge. pp. 25-78.
- Commons, M.L., Woodford, M., Trudeau, E.J. (1991). "How Each Reinforcer Contributes to Value: 'Noise' Must Reduce Reinforcer Value Hyperbolically". In: Commons, M.L., Nevin, J.A., Davison, M.C. (Eds.), *Signal Detection: Mechanisms, Models, and Applications*. Lawrence Erlbaum Associates, Hillsdale. pp. 139-168.
- Cropper, M. L., & Laibson, D. I. (1998). *The implications of hyperbolic discounting for project evaluation*. World Bank Publications.
- Cubitt, R.P., Read, D. (2007). "Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption?" *Experimental Economics* 10 (4), 369–89. doi:10.1007/s10683-006-9140-2.
- Currie, J. (2004). *The Take Up of Social Benefits*. National Bureau of Economic Research Working Paper No. 10488.
- Dai, X. and Fishbach, A., (2013). "When waiting to choose increases patience." *Organizational Behavior and Human Decision Processes*, 121(2), pp.256-266.
- Dasgupta, P., Maskin, E. (2005). "Uncertainty and Hyperbolic Discounting". *American Economic Review* 95 (4), 1290-1299.
- Dean, M., Sautmann, A. (2018). *Credit Constraints and the Measurement of Time Preferences*. Working Paper.
- Deck, C., Jahedi, S. (2015). "The Effect of Cognitive Load on Economic Decision Making: A Survey and New Experiments." *European Economic Review* 78, 97–119. doi:10.1016/j.eurocorev.2015.05.004.
- Dekel, E., Lipman, B.L., Rustichini, A. (2001). "Representing Preferences with a Unique Subjective State Space". *Econometrica* 69 (4), 891-934.
- Dekel, E., Lipman, B.L., Rustichini, A. (2009). "Temptation-Driven Preferences." *The Review of Economic Studies* 76 (3), 937–71. doi:10.1111/j.1467-937X.2009.00560.x.
- DellaVigna, S., Malmendier, U. (2004). "Contract Design and Self-Control: Theory and Evidence." *The Quarterly Journal of Economics* 119 (2), 353–402.
- DellaVigna, S., Malmendier, U. (2006). "Paying Not to Go to the Gym." *The American Economic Review* 96 (3), 694–719.
- DellaVigna, S., Paserman, M.D. (2005). "Job Search and Impatience." *Journal of Labor Economics* 23 (3), 527–588.
- DellaVigna, S, A. Lindner, B. Reizer, and J. Schmieder. (2017). "Reference-dependent Job Search: Evidence from Hungary." *The Quarterly Journal of Economics* 132(4), 1969-2018.
- DeSteno, D., Li, Y., Dickens, L., Lerner, J.S. (2014). "Gratitude: A Tool for Reducing Economic Impatience". *Psychological Science* 25 (6), 1262-1267.

- Dohmen, T., Falk, A., Huffman, D., Sunde, U. (2010). "Are Risk Aversion and Impatience Related to Cognitive Ability?" *American Economic Review* 100 (3), 1238-1260.
- Dreyfus, M.K., Viscusi, W.K. (1995). "Rates of Time Preference and Consumer Valuations of Automobile Safety and Fuel Efficiency." *Journal of Law and Economics* 38 (1), 79–105.
- Duckworth, A., Peterson, C., Matthews, M. D., Kelly, D.R. (2007). "Grit: Perseverance and Passion for Long-Term Goals". *Journal of Personality and Social Psychology* 92, 1087-1101.
- Duckworth, A., Gross, J.J. (2014). "Self-Control and Grit: Related but Separable Determinants of Success". *Current Directions in Psychological Science* 23 (5), 319-325.
- Duckworth, A.L., Laibson, D., Milkman, K.L. Forthcoming. "Beyond Willpower: Strategies for Reducing Failures of Self-Control." *Psychological Science in the Public Interest*.
- Duflo, E., Kremer, M., Robinson, J. (2011). "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya". *American Economic Review* 101 (6), 2350-2390.
- Dushi, I., Munell, A.H., Sanzenbacher, G.T., Webb, A. (2015). "Do Households Increase Their Savings When the Kids Leave Home?" Center for Retirement Research at Boston College Working Paper No. 2015-26.
- Dweck, C. S. (1986). "Motivational processes affecting learning". *American Psychologist* 41, 1040–1048.
- Ebert, J.E.J., Prelec, D. (2007). "The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future." *Management Science* 53 (9), 1423–38. doi:10.1287/mnsc.1060.0671.
- Ericson, K.M. (2011). "Forgetting We Forget: Overconfidence and Memory". *Journal of the European Economic Association* 9 (1), 43-60.
- Ericson, K.M. (2017). "On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation." *Journal of the European Economic Association* 15 (3), 692-719.
- Ericson, K.M., White, J.M., Laibson, D., Cohen, J.D. (2015). "Money Earlier or Later? Simple Heuristics Explain Intertemporal Choices Better than Delay Discounting Does." *Psychological Science* 26 (6), 826-833. doi: 10.1177/0956797615572232.
- Ericson, K.M., Noor, J. (2015). Delay Functions as the Foundation of Time Preference: Testing for Separable Discounted Utility. National Bureau of Economic Research Working Paper No. 21095.
- Epper, T., Fehr-Duda, H. (2015a). "The Missing Link: Unifying Risk Taking and Time Discounting". Working Paper No. 096, University of Zurich.
- Epper, T., Fehr-Duda, H. (2015b). "Comment on 'Risk Preferences are Not Time Preferences': Balancing on a Budget Line." *American Economic Review* 105 (7), 2261-2271.
- Epper, T., Fehr-Duda, H., Bruhin, A. (2011). "Viewing the Future through a Warped Lens: Why Uncertainty Generates Hyperbolic Discounting." *Journal of Risk and Uncertainty* 43 (3), 169–203. doi:10.1007/s11166-011-9129-x.
- Epstein, L.G., Zin, S.E. (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4), 937–969.
- Estle, S.J., Green, L., Myerson, J., Holt, D.D. (2007). "Discounting of Monetary and Directly Consumable Rewards." *Psychological Science* 18 (1), 58–63. doi:10.1111/j.1467-9280.2007.01849.x.
- Fadlon, I., Laibson, D. (2017). Paternalism and Pseudo-Rationality. National Bureau of Economic Research Working Paper No. 23620.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., Sunde, U. (2015). "The Nature and Predictive Power of Preferences: Global Evidence." IZA Discussion Papers No. 9504.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., Sunde, U. (2016). "The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences" IZA Discussion Papers No. 9674.
- Farkas, S., Johnson, J. (1997). *Miles to Go: A Status Report on Americans' Plans for Retirement*. Public Agenda, New York.
- Fedyk, A., (2018). "Asymmetric Naivete: Beliefs about Self-Control". SSRN Scholarly Paper ID 2727499, Social Science Research Network.

- Feldstein, M. (1985). "The Optimal Level of Social Security Benefits". *The Quarterly Journal of Economics* 100 (2), 303-320.
- Fernández-Villaverde, J., Mukherji, A. (2006). "Can We Really Observe Hyperbolic Discounting?". Working Paper.
- Fischer, C. (2001). "Read this paper later: procrastination with time-consistent preferences". *Journal of Economic Behavior & Organization* 46 (3), 249-269.
- Fox, M.V. (1983). "Ancient Egyptian Rhetoric". *Rhetorica: A Journal of the History of Rhetoric* 1 (1), 9-22.
- Frakes, M.D., Wasserman, M.F. (2017). Procrastination in the Workplace: Evidence from the U.S. Patent Office. National Bureau of Economic Research Working Paper No. 22987.
- Frederick, S. (2005). "Cognitive Reflection and Decision Making." *The Journal of Economic Perspectives* 19 (4), 25-42.
- Frederick, S., Loewenstein, G. (2008). "Conflicting motives in evaluations of sequences". *Journal of Risk and Uncertainty* 37 (2-3), 221-235.
- Frederick, S., Loewenstein, G., O'Donoghue, T. (2002). "Time discounting and time preference: A critical review". *Journal of Economic Literature* 40 (2), 351-401.
- Freeman, D., Manzini, P., Mariotti, M., Mittone, L. (2016). "Procedures for Eliciting Time Preferences." *Journal of Economic Behavior and Organization* 126, 235-242. <https://doi.org/10.1016/j.jebo.2016.03.017>
- Fudenberg, D., Levine, D.K. (2006). "A Dual-Self Model of Impulse Control." *The American Economic Review* 96 (5), 1449-1476.
- Fudenberg, D., Levine, D.K. (2011). "Risk, Delay, and Convex Self-Control Costs." *American Economic Journal: Microeconomics* 3 (3), 34-68. doi:10.1257/mic.3.3.34.
- Fudenberg, D., Levine, D.K. (2012). "Timing and self-control". *Econometrica* 80 (1), 1-42.
- Gabaix, X. (2014). "A Sparsity-Based Model of Bounded Rationality". *The Quarterly Journal of Economics* 129 (4), 1661-1710.
- Gabaix, X., Laibson, D. (2017). Myopia and Discounting. Working Paper.
- Ganong, P., Noel, P. (2017). Consumer Spending During Unemployment: Positive and Normative Implications. Working Paper.
- 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.
- Glimcher, P.W., Kable, J., Louie, K. (2007). "Neuroeconomic Studies of Impulsivity: Now or Just as Soon as Possible?" *The American Economic Review* 97 (2), 142-147.
- Goda, G.S., Levy, M.R., Manchester, C.F., Sojourner, A., Tasoff, J. (2018). "Mechanisms Behind Retirement Saving Behavior: Evidence from Administrative and Survey Data". TIAA Institute Research Dialogue 140.
- Gollier, C. (2013). "The Debate on Discounting: Reconciling Positivists and Ethicists." *Chicago Journal of International Law* 13 (2), 549-564.
- Golsteyn, B.H.H., Grönqvist, H., Lindahl, L. (2013). "Time Preferences and Lifetime Outcomes". IZA Discussion Paper No. 7165.
- Goulder, L.H., Williams, R.C. (2012). "The Choice of Discount Rate for Climate Change Policy Evaluation." *Climate Change Economics* 03 (04).
- Gross, D.B., Souleles, N.S. (2002). "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *The Quarterly Journal of Economics* 117 (1), 149-185.
- Gross, M.J., Notowidigdo, J.W. (2016). The Marginal Propensity to Consume Over the Business Cycle. National Bureau of Economic Research Working Paper No. 22518.
- Gul, F., Pesendorfer, W. (2001). "Temptation and Self-Control." *Econometrica* 69 (6), 1403-1435. doi:10.1111/1468-0262.00252.
- Gul, F., Pesendorfer, W. (2004). "Self-Control and the Theory of Consumption". *Econometrica* 72 (1), 119-158.
- Gustman, A.L., Steinmeier, T.L. (2012). "Policy Effects in Hyperbolic vs. Exponential Models of Consumption and Retirement." *Journal of Public Economics* 96 (5-6), 465-73. doi:10.1016/j.jpubeco.2012.02.001.

- Haider, S.J., Stephens, M., Jr. (2007). "Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations". *The Review of Economics and Statistics* 89 (2), 247-264.
- Halevy, Y. (2005). *Diminishing Impatience: Disentangling Time Preference from Uncertain Lifetime*. Working Paper, University of British Columbia.
- Halevy, Y. (2008). "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect". *American Economic Review* 98 (3), 1145-1162.
- Halevy, Y. (2014). "Some Comments on the Use of Monetary and Primary Rewards in the Measurement of Time Preferences". Technical Report, University of British Columbia.
- Halevy, Y. (2015). "Time Consistency: Stationarity and Time Invariance". *Econometrica* 83 (1), 335-352.
- Hamermesh, D.S. (1984). "Consumption During Retirement: The Missing Link in the Life Cycle". *The Review of Economics and Statistics* 66 (1), 1-7.
- Hardisty, D.J., Thompson, K.F., Krantz, D.H., Weber, E.U. (2013). "How to Measure Time Preferences: An Experimental Comparison of Three Methods." *Judgment and Decision Making* 8 (3), 236–49.
- Hare, T.A., Camerer, C.F., Rangel, A. "Self-control in decision-making involves modulation of the vmPFC valuation system". *Science* 324 (5927), 646-648.
- Harris, C., Laibson, D. (2001). "Dynamic choices of hyperbolic consumers". *Econometrica* 69 (4), 935-957. doi:10.1111/1468-0262.00225.
- Harris, C.J., Laibson, D. (2002). "Hyperbolic Discounting and Consumption." In: Dewatripont, M., Hansen, L.P., Turnovsky, S. (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Volume 1, Eighth World Congress*. Cambridge University Press, Cambridge.
- Harris, C., Laibson, D. (2012). "Instantaneous Gratification". *The Quarterly Journal of Economics* 128 (1), 205–248.
- Harrison, G.W., Lau, M.I., Williams, M.B. (2002). "Estimating Individual Discount Rates in Denmark: A Field Experiment." *American Economic Review* 92 (5), 1606–1617.
- Haushofer, J., Fehr, E. (2014). "On the psychology of poverty". *Science* 344 (6186), 862-867.
- Hausman, J.A. (1979). "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *The Bell Journal of Economics* 10 (1), 33–54.
- Heckman, J., Pinto, R., Savelyev, P. (2013). "Understanding the Mechanism through Which an Influential Early Childhood Program Boosted Adult Outcomes". *American Economic Review* 103 (6), 2052-2086.
- Heckman, J.J., Stixrud, J., Urzua, S. (2006). "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior". *Journal of Labor Economics* 24 (3), 411-482.
- Heidhues, P., Kosegi, B. (2010). "Exploiting Naïvete about Self-Control in the Credit Market". *American Economic Review* 100 (5), 2279-2303.
- Herrnstein, R.J. (1961). "Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Behavior* 4 (3), 267–272. doi:10.1901/jeab.1961.4-267.
- Hinson, J.M., Jameson, T.L., Whitney, P. (2003). "Impulsive Decision Making and Working Memory". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29 (2), 298-306.
- Hoch, S.J., Loewenstein, G.F. (1991). "Time-inconsistent preferences and consumer self-control". *Journal of Consumer Research* 17 (4), 492-507.
- Holcomb, J.H., Nelson, P.S. (1992). "Another Experimental Look at Individual Time Preference." *Rationality and Society* 4 (2), 199–220. doi:10.1177/1043463192004002006.
- Houser, D., Schunk, D., Winter, J., Xiao, E. (2018). "Temptation and Commitment in the Laboratory." *Games and Economic Behavior* 107, 329-344.
- Hurd, M.D., Rohwedder, S. (2013). "Heterogeneity in spending change at retirement". *The Journal of the Economics of Ageing* 1-2, 60-71.
- Imas, A., Kuhn, M.A., Mironova, V. (2017). "Waiting to Choose". Working Paper, Carnegie Mellon University.
- Jachimowicz, J., Duncan, S., Weber, E. U., and Johnson, E. J., (2018). "When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects" SSRN working paper.

- Jackson, M. Yariv, L. (2014). "Present Bias and Collective Dynamic Choice in the Lab." *American Economic Review* 104 (12), 4184-4204.
- Jackson, M. Yariv, L. (2015). "Collective Dynamic Choice: The Necessity of Time Inconsistency." *American Economic Journal: Microeconomics* 7 (4), 150-178.
- Jéhiel, P. (1995). "Limited Horizon Forecast in Repeated Alternate Games." *Journal of Economic Theory* 67 (2), 497-519.
- Johnson, D.S., Parker, J.A., Souleles, N.S. (2006). "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5), 1589-1610.
- Johnson, E.J., Goldstein, D. (2003). "Do Defaults Save Lives?". *Science* 302 (5649), 1338-1339.
- Kable, J.W., Glimcher, P.W. (2010). "An "as soon as possible" effect in human intertemporal decision making: behavioral evidence and neural mechanisms". *Journal of Neurophysiology* 103 (5), 2513-2531.
- Kahneman, D., Deaton, A. (2010) "High income improves evaluation of life but not emotional well-being". *Proceedings of the National Academy of Sciences* 107 (38), 16489-16493.
- Kahneman D., Fredrickson B.L., Schreiber C.A., Redelmeier D.A. (1993). "When more pain is preferred to less: Adding a better end". *Psychological Science* 4 (6), 401-405.
- Kahneman, D., Tversky, A. (1979). "Prospect Theory: An Analysis of Decision Under Risk". *Econometrica* 47 (2), 263-291.
- Kaplan, G., Violante, G.L. (2014). "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4), 1199-1239.
- Kaplan, G., Violante, G.L., Weidner, J. (2014). The wealthy hand-to-mouth. National Bureau of Economic Research Working Paper No. 20073.
- Kaur, S., Kremer, M., Mullainathan, S. (2010). "Self-control and the development of work arrangements". *American Economic Review* 100 (2), 624-628.
- Kaur, S., Kremer, M., Mullainathan, S. (2015). "Self-control at work". *Journal of Political Economy* 123 (6), 1227-1277.
- Keren, G., Roelofsma, P. (1995). "Immediacy and Certainty in Intertemporal Choice". *Organizational Behavior and Human Decision Processes* 63 (3), 287-297.
- Kirby, K.N. (1997). "Bidding on the Future: Evidence against Normative Discounting of Delayed Rewards." *Journal of Experimental Psychology: General* 126 (1), 54-70. doi:10.1037/0096-3445.126.1.54.
- Kirby, K.N. (2009). "One-Year Temporal Stability of Delay-Discount Rates." *Psychonomic Bulletin & Review* 16 (3), 457-62. doi:10.3758/PBR.16.3.457.
- Kirby, K.N., Herrnstein, R.J. (1995). "Preference Reversals Due to Myopic Discounting of Delayed Reward." *Psychological Science* 6 (2), 83-89. doi:10.1111/j.1467-9280.1995.tb00311.x.
- Koopmans, T. (1960). "Stationary ordinal utility and impatience." *Econometrica*, 28 pp. 287-309.
- Kőszegi, B., Rabin, M. (2009). "Reference-dependent consumption plans". *American Economic Review* 99 (3), 909-936.
- Kőszegi, B., Szeidl, A. (2013). "A model of focusing in economic choice". *The Quarterly Journal of Economics* 128 (1), 53-104.
- Kreps, D. M. (1979). "A Representation Theorem for "Preference for Flexibility"". *Econometrica* 47 (3), 565-577.
- Kreps, D.M. (1998). "Anticipated Utility and Dynamic Choice". In: Jacobs, D., Kalai, E., Kamien, M. (Eds.), *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*, Econometric Society Monographs. Cambridge University Press, Cambridge.
- Kreps, D. M., & Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica: journal of the Econometric Society*, 185-200.
- Krupka, E.L., Stephens, M., Jr. (2013). "The Stability of Measured Time Preferences." *Journal of Economic Behavior & Organization* 85, 11-19. doi:10.1016/j.jebo.2012.10.010.
- Krusell, P., Kuruşçu, B., Smith, A.A., Jr. (2010). "Temptation and taxation". *Econometrica* 78 (6), 2063-2084. doi:10.3982/ECTA8611.

- Kuchler, T. (2015). Sticking to Your Plan: Hyperbolic Discounting and Credit Card Debt Paydown. Working Paper.
- Kuchler, T., Pagel, M. (2018). Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown. National Bureau of Economic Research Working Paper No. 24881.
- Laibson, D. (1997). "Golden Eggs and Hyperbolic Discounting". *The Quarterly Journal of Economics* 112 (2), 443-477. doi:10.1162/003355397555253.
- Laibson, D. (2001). "A Cue-Theory of Consumption." *The Quarterly Journal of Economics* 116 (1), 81–119.
- Laibson, D. (2015). "Why Don't Present-Biased Agents Make Commitments?". *American Economic Review* 105 (5), 267-272.
- Laibson, D. (2018). "Private Paternalism, the Commitment Puzzle, and Model-Free Equilibrium." *AEA Papers and Proceedings* 108, 1-21.
- Laibson, D., Maxted, P., Repetto, A., Tobacman, J. (2017). Estimating Discount Functions with Consumption Choices over the Lifecycle. Working Paper.
- Laibson, D., Repetto, A., Tobacman, J. (2003). "A Debt Puzzle". In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press, Princeton.
- Laibson, D., Repetto, A., Tobacman, J., (1998). "Self-Control and Saving for Retirement". *Brookings Papers on Economic Activity* 1998 (1), 91-196.
- Lavecchia, A., Liu, H., Oreopoulos, P. (2014). Behavioral Economics of Education: Progress and Possibilities. National Bureau of Economic Research Working Paper No. 20609.
- Lawrance, E.C. (1991). "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* 99 (1), 54–77.
- Lawyer, S.R., Williams, S.A., Prihodova, T., Rollins, J.D., Lester, A.C. (2010). "Probability and Delay Discounting of Hypothetical Sexual Outcomes." *Behavioural Processes* 84 (3), 687–92. doi:10.1016/j.beproc.2010.04.002.
- Lerner, J.S., Li, Y., Weber, E.U. (2013). "The Financial Cost of Sadness". *Psychological Science* 24 (1), 72-79.
- Lipman, B.L., Pesendorfer, W. (2013). "Temptation". In: Acemoglu, D., Arellano, M., Dekel, E. (Eds.), *Advances in Economics and Econometrics: Tenth World Congress Volume 1*. Cambridge University Press, New York.
- Lockwood, B. (2016). "Optimal Income Taxation with Present Bias". Working Paper.
- Loewenstein, G. (1987). "Anticipation and the Valuation of Delayed Consumption." *The Economic Journal* 97 (387), 666-684. doi:10.2307/2232929.
- Loewenstein, G. (1988). "Frames of Mind in Intertemporal Choice." *Management Science* 34 (2), 200–214. doi:10.1287/mnsc.34.2.200.
- Loewenstein, G. (1992). "The fall and rise of psychological explanations in the economics of intertemporal choice". In: Loewenstein G., Elster J. (Eds.), *Choice over time*. Russell Sage Foundation, New York.
- Loewenstein, G. (1996). "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes* 65 (3), 272–92. doi:10.1006/obhd.1996.0028.
- Loewenstein, G., Elster, J. (Eds.) (1992). *Choice over time*. Russell Sage Foundation, New York.
- Loewenstein, G., O'Donoghue, T. (2004). *Animal Spirits: Affective and Deliberative Processes in Economic Behavior*. SSRN Scholarly Paper ID 539843, Social Science Research Network.
- Loewenstein, G., Prelec, D. (1992). "Anomalies in intertemporal choice: Evidence and an interpretation". *The Quarterly Journal of Economics* 107 (2), 573-597.
- Loewenstein, G.F., Prelec, D. (1993). "Preferences for sequences of outcomes". *Psychological Review* 100 (1), 91-108.
- Loewenstein, G., Read, D., Baumeister, R.F. (Eds.). (2003). *Time and Decision*. Russell Sage Foundation, New York.
- Loewenstein, G., Sicherman, N. (1991). "Do workers prefer increasing wage profiles?". *Journal of Labor Economics* 9 (1), 67-84.
- Madrian, B.C., Shea, D.F. (2001). "The power of suggestion: Inertia in 401 (k) participation and savings behavior". *The Quarterly Journal of Economics* 116 (4), 1149-1187.

- Maital, S., Maital, S. (1977). "Time Preference, Delay of Gratification and the Intergenerational Transmission of Economic Inequality: A Behavioral Theory of Income Distribution." In: Ashenfelter, O.C., Oates, W.E. (Eds.), *Essays in Labor Market Analysis: In Memory of Yochanan Peter Comay*. John Wiley and Sons, New York.
- Mariger, R.P. (1987). "A Life-cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results". *Econometrics* 55 (3), 533-557.
- Marshall, A. (1891). *Principles of Economics*, 2nd ed. Macmillan, London.
- Mastrobuoni, G., Weinberg, M. (2009). "Heterogeneity in Intra-Monthly Consumption Patterns, Self-Control, and Savings at Retirement." *American Economic Journal: Economic Policy* 1 (2), 163–189. doi:10.1257/pol.1.2.163.
- McClure, S.M., Ericson, K.M., Laibson, D.I., Loewenstein, G., Cohen, J.D. (2007). "Time Discounting for Primary Rewards." *The Journal of Neuroscience* 27 (21), 5796–5804. doi:10.1523/JNEUROSCI.4246-06.2007.
- 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. doi:10.1126/science.1100907.
- McGuire, J.T., Kable, J.W. (2012). "Decision Makers Calibrate Behavioral Persistence on the Basis of Time-Interval Experience". *Cognition* 124 (2), 216-226.
- McGuire, J.T., Kable, J.W. (2013). "Rational Temporal Predictions Can Underlie Apparent Failures to Delay Gratification". *Psychological Review* 120 (2), 395-410.
- McIntosh, D. (1969). *The Foundations of Human Society*. University of Chicago Press, Chicago.
- Meier, S., Sprenger, C. (2010). "Present-Biased Preferences and Credit Card Borrowing." *American Economic Journal: Applied Economics* 2 (1), 193–210. doi:10.1257/app.2.1.193.
- Meier, S., Sprenger, C.D. (2011). "Time Discounting Predicts Creditworthiness." *Psychological Science* 23 (1) 56-58. doi:10.1177/0956797611425931.
- Meier, S., Sprenger, C. (2015). "Temporal Stability of Time Preferences." *Review of Economics and Statistics* 97 (2), 273-286.
- 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 (26), 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 (1), 92-93.
- Milkman, K.L., Rogers, T., Bazerman, M.H. (2009). "Highbrow films gather dust: Time-inconsistent preferences and online DVD rentals". *Management Science* 55 (6), 1047-1059.
- Mischel, W., Shoda, Y., Rodriguez M.I. (1989). "Delay of Gratification in Children." *Science* 244 (4907): 933–938. doi:10.1126/science.2658056.
- Mischel, W., Ebbesen, E.B. (1970). "Attention in Delay of Gratification." *Journal of Personality and Social Psychology* 16 (2), 329–337. doi:10.1037/h0029815.
- Miao, B., Zhong, S. (2015). "Comment on 'Risk Preferences are Not Time Preferences': Separating Risk and Time Preference." *American Economic Review* 105 (7), 2272-2286.
- Moser, C., & Olea de Souza e Silva, P. (2017). "Optimal paternalistic savings policies."
- Moffitt, T.E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R.J., Harrington, H., ... & Sears, M. R. (2011). "A gradient of childhood self-control predicts health, wealth, and public safety". *Proceedings of the National Academy of Sciences* 108 (7), 2693-2698.
- Muraven, M., Baumeister, R.F. (2000). "Self-regulation and depletion of limited resources: does self-control resemble a muscle?". *Psychological Bulletin* 126 (2), 247-259.
- Muraven, M., Tice, D.M., Baumeister, R.F. (1998). "Self-Control as a Limited Resource: Regulatory Depletion Patterns." *Journal of Personality and Social Psychology* 74 (3), 774–789. doi:10.1037/0022-3514.74.3.774.
- Noor, J. (2007). "Commitment and Self-control." *Journal of Economic Theory*, 135(1), 1-34.
- Noor, J. (2009). "Hyperbolic Discounting and the Standard Model: Eliciting Discount Functions." *Journal of Economic Theory* 144 (5), 2077–2083.

- Noor, J. (2011). "Temptation and revealed preference." *Econometrica*, 79(2), 601-644.
- Nussbaum, M.C. (2001). "'This story isn't true': Madness, reason and recantation in the Phaedrus". In: *The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philosophy*, 200-234. Cambridge University Press, Cambridge.
- 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". *The Quarterly Journal of Economics* 114 (3), 769-816.
- O'Donoghue, T., Rabin, M. (2001). "Choice and Procrastination". *The Quarterly Journal of Economics* 116 (1), 121-160.
- Odum, A.L., Rainaud, C.P. (2003). "Discounting of Delayed Hypothetical Money, Alcohol, and Food." *Behavioural Processes* 64 (3), 305-313. doi:10.1016/S0376-6357(03)00145-1.
- Odum, A.L., Baumann, A.A.L., Rimington, D.D. (2006). "Discounting of Delayed Hypothetical Money and Food: Effects of Amount." *Behavioural Processes* 73 (3), 278-84. doi:10.1016/j.beproc.2006.06.008.
- Olafsson, A., Pagel, M. (2018). *The Retirement-Consumption Puzzle: New Evidence from Personal Finances*. National Bureau of Economic Research Working Paper No. 24405.
- Olea, J.L.M., Strzalecki, T. (2014). "Axiomatization and Measurement of Quasi-Hyperbolic Discounting." *The Quarterly Journal of Economics* 129 (3) 1449-1499.
- Oster, S. M., & Morton, F. M. S. (2005). Behavioral biases meet the market: The case of magazine subscription prices. *The BE Journal of Economic Analysis & Policy*, 5(1).
- Pallais, A. (2015). "Small Difference That Matter: Mistakes in Applying to College". *Journal of Labor Economics* 33 (2), 493-520.
- Parker, J.A., Souleles, N.S., Johnson, D.S., McClelland, R. (2013). "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6), 2530-2553.
- Paserman, M.D. (2008). "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation." *The Economic Journal* 118 (531), 1418-52. doi:10.1111/j.1468-0297.2008.02175.x.
- Pew Charitable Trusts. (2012). "Who Borrows, Where They Borrow, and Why - Pew Center on the States". *Payday Lending in America Report Series*, Pew Charitable Trusts, Philadelphia.
- Phelps, E.S., Pollak, R.A. (1968). "On Second-Best National Saving and Game-Equilibrium Growth." *The Review of Economic Studies* 35 (2), 185-199. doi:10.2307/2296547.
- Pigou, A.C. (1920). *The Economics of Welfare*. Macmillan, London.
- Plato. (1925). *Plato in Twelve Volumes*. Translated by Harold N. Fowler, vol. 9. Cambridge, MA, Harvard University Press; London, William Heinemann Ltd.
- Prelec, D., Loewenstein, G. (1991). "Decision Making over Time and Under Uncertainty: A Common Approach". *Management Science* 37 (7), 770-786.
- Prelec, D., Loewenstein, G. (1998). "The red and the black: Mental accounting of savings and debt". *Marketing Science* 17 (1), 4-28.
- Quiggin, J., Horowitz, J. (1995). "Time and Risk". *Journal of Risk and Uncertainty* 10 (1), 37-55.
- Rabin, M. (2000). "Risk Aversion and Expected-utility Theory: A Calibration Theorem". *Econometrica* 68 (5), 1281-1292.
- Radcliffe, E.S. (2018). *Hume, Passion, and Action*. Oxford University Press, Oxford.
- Ramsey, F. P. (1928). A mathematical theory of saving." *Economic Journal* 38, 543-559.
- Rangel, A., Camerer, C., Montague, P.R. (2008). "A framework for studying the neurobiology of value-based decision making". *Nature reviews neuroscience* 9 (7), 545-556.
- Read, D. (2001). "Is Time-Discounting Hyperbolic or Subadditive?" *Journal of Risk and Uncertainty* 23 (1), 5-32. doi:10.1023/A:1011198414683.
- Read, D., Frederick, S., Scholten, M. (2013). "DRIFT: An Analysis of Outcome Framing in Intertemporal Choice". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 39 (2), 573-588.

- Read, D., Loewenstein, G., Kalyanaraman, S. (1999). "Mixing Virtue and Vice: Combining the Immediacy Effect and the Diversification Heuristic." *Journal of Behavioral Decision Making* 12 (4), 257–273.
- Read, D., Powell, M. (2002). "Reasons for sequence preferences". *Journal of Behavioral Decision Making* 15 (5), 433-460.
- Read, D., Roelofsma, P.H.M.P. (2003). "Subadditive versus Hyperbolic Discounting: A Comparison of Choice and Matching." *Organizational Behavior and Human Decision Processes* 91 (2), 140–153.
- 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.
- Reuben, E., Sapienza, P., Zingales, L. (2010). "Time Discounting for Primary and Monetary Rewards." *Economics Letters* 106 (2), 125–127.
- Reuben, E., Sapienza, P., Zingales, L. (2015). "Procrastination and Impatience." *Journal of Behavioral and Experimental Economics* 58, 63-76.
- Rick, S. I., Cryder, C. E., & Loewenstein, G. (2007). "Tightwads and Spendthrifts." *Journal of Consumer Research*, 34(6), 767-782.
- Roberts, B.W., Chernyshenko, O.S., Stark, S., Goldberg, L.R. (2005). "The Structure of Conscientiousness: An Empirical Investigation Based on Seven Major Personality Questionnaires". *Personnel Psychology* 58 (1), 103–139.
- Robson, A., Samuelson, L. (2009). "The Evolution of Time Preference with Aggregate Uncertainty." *American Economic Review* 99 (5), 1925-1953.
- Royer, H., Stehr, M., Sydnor, J. (2015). "Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company". *American Economic Journal: Applied Microeconomics* 7 (3), 51-84.
- Rubinstein, A. (2003). "Economics and Psychology? The Case of Hyperbolic Discounting". *International Economic Review* 44 (4), 1207-1216.
- Ryder, H.E., Heal, G. (1973). "Optimal Growth with Intertemporally Dependent Preferences". *Review of Economic Studies* 40 (1), 1-31.
- Sadoff, S., Samek, A., Sprenger, C. (2015). "Dynamic Inconsistency in Food Choice: Experimental Evidence from a Food Desert". Becker Friedman Institute for Research in Economics Working Paper No. 2572821, CESR-Schaeffer Working Paper No. 2015-027.
- Saito, K. (2011). "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect: Comment." *The American Economic Review* 101 (5), 2271-2275.
- Samuelson, P. (1937). "A Note on Measurement of Utility." *Review of Economic Studies* 4 (2), 155–161. doi:10.2307/2967612.
- Samuelson, W., Zeckhauser, R. (1988). "Status Quo Bias in Decision Making". *Journal of Risk and Uncertainty* 1, 7-59.
- Schelling, T.C. (1995). "Intergenerational Discounting." *Energy Policy* 23 (4–5), 395–401.
- Schilbach, F. (2017). "Alcohol and Self-Control: A Field Experiment in India." Working Paper.
- Schilbach, F., Schofield, H., Mullainathan, S. (2016). "The Psychological Lives of the Poor". *American Economic Review: Papers and Proceedings* 106 (5), 435-440.
- Scholten, M., Read, D. (2006). "Discounting by Intervals: A Generalized Model of Intertemporal Choice." *Management Science* 52 (9), 1424–1436.
- Scholz, J.K., Seshadri, A., Khitatrakun, S. (2006). "Are Americans Saving "Optimally" for Retirement?". *Journal of Political Economy* 114 (4), 607-643.
- Schultz, W., Dayan, P., Montague R.R. (1997). "A Neural Substrate of Prediction and Reward". *Science* 275, 1593-1599.
- Sen, A.K. (1957). "A Note on Tinbergen on the Optimum Rate of Saving." *The Economic Journal* 67 (268), 745-748.

- Shah, A. K., Mullainathan, S., Shafir, E. (2012). "Some consequences of having too little". *Science* 338 (6107), 682-685.
- Shapiro, J. M. (2005). "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *Journal of Public Economics* 89, 303–325.
- Shapiro, M.D., Slemrod, J. (2003). "Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys." *Tax Policy and the Economy* 17, 83-109.
- Shea, J. (1995). "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis." *American Economic Review* 85(1), 186–200.
- Shefrin, H.M., Thaler, R.H. (1988). "The Behavioral Life-cycle Hypothesis." *Economic Inquiry* 26 (4), 609-643.
- Shiv, B., Fedorikhin, A. (1999). "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making". *Journal of Consumer Research* 26 (3), 278-292.
- Shoda, Y., Mischel, W., Peake, P.K. (1990). "Predicting Adolescent Cognitive and Self-Regulatory Competencies From Preschool Delay of Gratification: Identifying Diagnostic Conditions". *Developmental Psychology* 26 (6), 978-986.
- Shu, S., Gneezy, A. (2010). "Procrastination of Enjoyable Experiences". *Journal of Marketing Research* 47, 933-944.
- Simon, C.J., Warner, J.T., Pleeter, S. (2014). "Discounting, Cognition, and Financial Awareness: New Evidence from a Change in the Military Retirement System". *Economic Inquiry* 53 (1), 318–334.
- Sisk, V.F., Burgoyne, A.P., Sun, J., Butler, J.L., Macnamara, B.N. (2018). "To What Extent and Under Which Circumstances are Growth Mind-sets Important to Academic Achievement? Two Meta-analyses". *Psychological Science* 29 (4), 549-571.
- Skiba, P.M. (2012). "Regulation of Payday Loans: Misguided?". *Washington and Lee Law Review* 69 (2), 1023-1049.
- Skiba, P. M., & Tobacman, J. (2008). Payday loans, uncertainty and discounting: explaining patterns of borrowing, repayment, and default. *Vanderbilt Law and Economics Research Paper No. 08-33*. Available at <http://dx.doi.org/10.2139/ssrn.1319751>
- Skinner, J. (2007). "Are You Sure You're Saving Enough for Retirement?". *Journal of Economic Perspectives* 21 (3), 59-80.
- Smith, A. (1759). *The Theory of Moral Sentiments*. A. Miller, A. Kincaid, and J. Bell, Strand and Edinburgh.
- Solomon, L.J., Rothblum, E.D. (1984). "Academic procrastination: Frequency and cognitive-behavioral correlates". *Journal of Counseling Psychology* 31 (4), 503-509.
- Sozou, P.D. (1998). "On Hyperbolic Discounting and Uncertain Hazard Rates". *Proceedings of the Royal Society B: Biological Sciences* 265 (1409), 2015-2020.
- Stango, V., Zinman, J. (2009). "Exponential Growth Bias and Household Finance". *The Journal of Finance* 64 (6), 2807-2849.
- Stavins, J. (2000). "Credit Card Borrowing, Delinquency, and Personal Bankruptcy." *New England Economic Review* 2000 (July/August), 15–30.
- Steele, C.M., Josephs, R.A. (1990). "Alcohol Myopia: Its Prized and Dangerous Effects". *American Psychologist* 45 (8), 921-933.
- Sugden, R. (2004). "The Opportunity Criterion: Consumer Sovereignty without the Assumption of Coherent Preferences". *American Economic Review* 94, 1014–1033.
- Stephens, M. (2003). "'3rd of the Month': Do Social Security Recipients Smooth Consumption between Checks?" *The American Economic Review* 93 (1), 406–422.
- Stephens, M., Jr., Toohey, D. (2018). *Changes in Nutrient Intake at Retirement*. National Bureau of Economic Research Working Paper No. 24621.
- Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press. doi:10.1017/CBO9780511817434.
- Stern, N. (2008). "The Economics of Climate Change". *American Economic Review* 98 (2), 1-37.

- Stevenson, B., Wolfers, J. (2013). "Subjective well-being and income: Is there any evidence of satiation?". *American Economic Review* 103 (3), 598-604.
- Strotz, R.H. (1955). "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies* 23 (3), 165–180.
- Thaler, R. H., Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press.
- Sutter, M., Kocher, M.G., Glätzle-Rüetzler, D., Trautmann, S.T. (2013). "Impatience and Uncertainty: Experimental Decisions Predict Adolescents' Field Behavior." *American Economic Review* 103 (1), 510–531.
- Takahashi, T. (2005). "Loss of self-control in intertemporal choice may be attributable to logarithmic time-perception". *Medical Hypotheses* 65 (4), 691-693.
- Taubinsky, D. (2014). *From Intentions to Actions: A Model and Experimental Evidence of Inattentive Choice*. Working Paper.
- Thakral, N, and Tô, L. (2017). "Daily Labor Supply and Adaptive Reference Points." Working Paper.
- Thaler, R. (1981). "Some Empirical Evidence on Dynamic Inconsistency." *Economics Letters* 8 (3), 201–207.
- Thaler, R.H., Shefrin, H.M. (1981). "An Economic Theory of Self-Control." *Journal of Political Economy* 89 (2), 392–406.
- Thornton, R.L. (2008). "The Demand for, and Impact of, Learning HIV Status." *American Economic Review* 98 (5), 1829-63.
- Toussaert, S. (2018). "Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment". *Econometrica* 86 (5), 859-889.
- Tsukayama, E., Duckworth, A.L. (2010). "Domain-Specific Temporal Discounting and Temptation." *Judgment and Decision Making* 5 (2), 72–82.
- Ubfal, D. (2016). "How General are Time Preferences? Eliciting Good-specific Discount Rates". *Journal of Development Economics* 118, 150-170.
- Viscusi, W.K., Moore, M.J. (1989). "Rates of Time Preference and Valuations of the Duration of Life." *Journal of Public Economics* 38 (3), 297–317.
- Warner, J.T., Pleeter, S. (2001). "The Personal Discount Rate: Evidence from Military Downsizing Programs." *The American Economic Review* 91 (1), 33–53.
- Watts, T.W., Duncan, G.J., Quan, H. (2018). "Revisiting the Marshmallow Test: A Conceptual Replication Investigating Links Between Early Delay of Gratification and Later Outcomes". *Psychological Science* 29 (7), 1159-1177.
- Weber, B.J., Chapman, G.B. (2005). "The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?". *Organizational Behavior and Human Decision Processes* 96 (2), 104-118.
- Weber, E.U., Johnson, E.J., Milch, K.F., Changwong, H., Brodscholl, J.C., Goldstein, D.G. (2007). "Asymmetric discounting in intertemporal choice: A query-theory account". *Psychological Science* 18 (6), 516-523.
- Weil, P. (1990). "Unexpected Utility in Macroeconomics." *The Quarterly Journal of Economics*, 105(1), 29-42.
- Weitzman, M. (2001). "Gamma Discounting". *American Economic Review* 91 (1), 260-271.
- Wertenbroch, K. (1998). "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice." *Marketing Science* 17 (4), 317–337.
- Wong, W. (2008). "How Much Time-inconsistency Is There and Does It Matter? Evidence of Self-awareness, Size, and Effects". *Journal of Economic Behavior & Organization* 68 (3-4), 645-656.
- Wulff, D.U., van den Bos, W. (2017). "Modelling Choices in Delay Discounting". *Psychological Science* 29 (11), 1890-1894.
- Yeager, D. S., Hanselman, P., Paunesku, D., Hulleman, C., Dweck, C., Muller, C., ... Duckworth, A. (2018). *Where and For Whom Can a Brief, Scalable Mindset Intervention Improve Adolescents' Educational Trajectories?* PsyArXiv Working Paper. doi:10.31234/osf.io/md2qa.

Zauberman, G., Kim, B. K., Malkoc, S. A., Bettman, J. R. (2009). "Discounting time and time discounting: Subjective time perception and intertemporal preferences". *Journal of Marketing Research* 46 (4), 543-556. doi:10.1509/jmkr.46.4.543.

Zinman, J. (2015). "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics* 7, 251–276.